NASA CR-170, 585

NASA-CR-170585
19840003749

A Reproduced Copy

OF

NASA CR-170, 585

Recommendations for NASA Research + Development in Artificial Intelligence (NASA-CR-170585)

Reproduced for NASA

by the

NASA Scientific and Technical Information Facility

FFNo 672 Aug 65

NF02259

BEST

AVAILABLE

COPY

170585

RECOMMENDATIONS FOR NASA RESEARCH AND DEVELOPMENT IN ARTIFICIAL INTELLIGENCE

Final Report

April 1983

By: David R. Brown
Advanced Computer Systems Department

Peter C. Cheeseman
Artificial Intelligence Center
Computer Science and Technology Division

Prepared for:

NASA/Goddard Space Flight Center
Greenbelt, Maryland 20771

Attention: Rodger A. Cliff, Code 402

Contract No. NAS5-27338
SRI Project 4716

SRI International
333 Ravenswood Avenue
Menlo Park, California 94025
(415) 326-6200
Cable: SRI INTL MPK
TWX: 910-373-2046

N84-11817

SRI International

RECOMMENDATIONS FOR NASA RESEARCH AND DEVELOPMENT IN ARTIFICIAL INTELLIGENCE

Final Report

April 1983

By: David R. Brown

Advanced Computer Systems Department

Peter C. Cheeseman

Artificial Intelligence Center
Computer Science and Technology Division

Prepared for:

NASA/Goddard Space Flight Center
Greenbelt, Maryland 20771

Attention: Rodger ... Cliff, Code 402

Contract No. NAS5-27338
SRI Project 4716

Approved:

Louis Fried, Director
Advanced Computer Systems Department

David H. Brandin, Vice President and Director
Computer Science and Technology Division

SRI International

This Page Intentionally Left Blank

ABSTRACT

Artificial intelligence, a branch of computer science, attempts to emulate human intelligence with computer systems based on symbol manipulation rather than numerical processing. The results of over two decades of artificial intelligence (AI) research are now beginning to be applied for natural-language access to computers, computer systems that act as expert consultants, and robotics. AI capabilities are also being developed for planning, design, image analysis and understanding, and other functions.

NASA, already aware of its need for AI in the years ahead, is beginning to develop its first applications. The potential for future applications is indeed enormous. AI will be used to improve some present-day functions and to make new programs and missions practicable. AI will be needed for the management of information, for engineering, and for administrative management. While the first applications will be in mission planning, AI will serve as a vitally important tool for diagnosis and repair of faults on future spacecraft. AI applications can be expected to increase dramatically in number and variety during the next five to ten years, in a manner analogous to the proliferation of computer applications in NASA during the last two decades. The space station should be planned from the start to incorporate AI techniques. An AI research group should be established in NASA, probably at Ames Research Center, with direct and continuing links to the AI research community, the entire range of applications in NASA, and NASA management. Although experienced AI specialists are scarce, the proposed AI research group could help train NASA personnel already skilled in system development, so that they could develop the artificial intelligence applications that will be needed by NASA.

This Page Intentionally Left Blank

CONTENTS

ILLUSTRATIONS

This Page Intentionally Left Blank

TABLES

PRECEDING PAGE BLANK NOT FILMED

This Page Intentionally Left Blank

GLOSSARY

ACRONYM - A model-based vision system.

AI - Artificial Intelligence. The capability of a machine to imitate intelligent human behavior.

AP - Automatic Programming. A branch of AI.

ARAMIS - Automation, Robotics, and Machine Intelligence Systems. An acronym used in an MIT study for the George C. Marshall Space Flight Center.

ARPANET - ARPA Network. A high-speed (50,000 baud), packet-switched, digital network developed by the Defense Advanced Research Projects Agency.

ASEE - American Society for Engineering Education, Washington, D.C.

AXAF - Advanced X-ray Astrophysics Facility. An x-ray telescope spacecraft.

CAD - Computer-Aided Design.

CAM - Computer-Aided Manufacturing.

CASE - Computer-Aided Systems Engineering. A system discussed at the 1981 NASA/ASEE Summer Study.

CAT - Computer-Aided Testing.

CETI - Communication with Extraterrestrial Intelligence.

CHI - A knowledge-based programming environment.

CMS - Command Management System, Goddard Space Flight Center.

CMU - Carnegie-Mellon University, Pittsburgh, Pennsylvania.

DBMS - Data Base Management System.

DEC - Digital Equipment Corporation, Maynard, Massachusetts.

DENDRAL - An expert system for inferring chemical structures.

DEVISER - An expert system for planning Voyager missions.

EL - An expert system for analysis of electrical circuits.

EXPERT - An expert system for medical consultation.

GE - General Electric Company, Fairfield, Connecticut.

GOES - Geostationary Operational Environmental Satellite.

GPS - General Problem Solver, an early problem solver using means-ends analysis.

GSFC - Goddard Space Flight Center, Greenbelt, Maryland.

GSP - Geostationary Platform, a proposed communications relay satellite.

HASP/SIAP - An expert system for signal processing.

IBM - International Business Machines, Armonk, New York.

ICAM - Integrated Computer-Aided Manufacturing. A program sponsored by the U.S. Air Force.

IESIS - Intelligent Earth-Sensing Information System. A hypothetical system described at the 1980 NASA/ASEE Summer Study.

INDUCE - An expert system for rule-guided inductive inference.

INTERLISP - Interactive LISP. A dialect of LISP.

IPAD - Integrated Programs for Aerospace Vehicle Design. A program sponsored by NASA Langley Research Center.

IPF - Image-Processing Facility, Goddard Space Flight Center.

JPL - Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California.

JSC - Lyndon B. Johnson Space Center, Houston, Texas.

KNEECAP - Knowledge-Based English-Enquiry Crew Activity Planning. A planning system based on KNOBS.

KNOBS - Knowledge-Based System An expert system for planning tactical air missions.

LANDSAT - A NASA satellite for observing earth resources.

LISP - List-Processing language. A programming language widely used in the AI community.

LSI - Large-Scale Integrated semiconductor circuits.

MACSYMA - An expert system for symbolic mathematics.

MIT - Massachusetts Institute of Technology, Cambridge, Massachusetts.

MITRE - The Mitre Corporation, Bedford, Massachusetts.

MOLGEN - An expert system for planning experiments in molecular genetics.

MSOCC - Multisatellite Operations Control Center, Goddard Space Flight Center.

MYCIN - An expert system for diagnosis and treatment of infectious diseases.

NASA - National Aeronautics and Space Administration, Washington, D.C.

NASTRAN - NASA Structural Analysis program.

NBS - National Bureau of Standards, Gaithersburg, Maryland.

NCC - Network Control Center, Goddard Space Flight Center.

NEEDS - NASA End-to-End Data System.

POCC - Project Operations Control Center, Goddard Space Flight Center.

PROLOG - Programming in Logic. An AI programming language.

PROSPECTOR - An expert system for mineral exploration.

RECON - A computer-based, bibliographic, NASA data base.

R1 - An expert system for configuring DEC VAX computers.

SETI - Search for Extraterrestrial Intelligence.

SIPE - System for Interactive Planning and Execution. An expert planning system developed at SRI International.

SP - Space Platform. A versatile platform for scientific and space applications research.

SRI - SRI International, Menlo Park, California.

STRIPS - A planning system for AI.

TMS - Teleoperator Maneuvering System. A multipurpose free-flying satellite tender.

VAX - Virtual-Address Extension. A 32-bit DEC computer with a virtual-memory operating system.

VLSI - Very-Large-Scale Integrated semiconductor circuits

ACKNOWLEDGMENT

During the preparation of this report, the authors benefited from their close association with the NASA technical officer, Rodger A. Cliff, an advantage made possible by his assignment to nearby Stanford University. His patience and understanding as the report gradually took form are greatly appreciated.

I INTRODUCTION

The purpose of this report is to make some reasonable and well-founded recommendations to NASA regarding the emerging new technology called artificial intelligence (AI). Although research in AI was started as a branch of computer science approximately at the time that NASA itself was created, AI is just beginning to be utilized in the real world. Its initial applications, though few in number, have generated considerable interest in AI's promising and multifaceted potential. This report provides a tutorial description of AI, takes a look at some of its possible applications in NASA, and concludes with recommendations for a comprehensive artificial intelligence program in NASA.

AI can be defined as a branch of computer science whose goal is to formalize the characteristics of intelligent behavior and to design, build, and comprehend machines that reason, plan, and perceive. These machines, i.e., computer systems, will have many of the cognitive skills associated with human beings and could therefore be used pervasively over a broad range of human intellectual activities, including all the sciences, the professions, industry, literature, and so on.

Fundamental research in AI is concerned primarily with the problems of knowledge representation and commonsense reasoning, and with the development of techniques for finding solutions (called heuristic search). Categories of AI systems include expert systems, planning, theorem proving, vision, natural language, robotics, automatic programming, and learning.

AI has potential for useful applications throughout NASA's activities. These include the management of information, engineering, institutional management, and the performance of previously impractical missions. Once initiated and properly implemented, AI is expected to proliferate in NASA the way computers did during the last twenty years. The following are examples of potential applications:

- An expert system for mission planning. Such an application, called DEVISER, is already being developed at JPL (Vere, 1981); it can be used at other NASA locations as well.
- An expert system to detect and correct faults on spacecraft.
- A vision system to analyze LANDSAT imagery.
- Intelligence for autonomous spacecraft in planetary exploration or in deep space.
- Control of high-performance future aircraft that humans cannot control alone.
- Automatic generation of computer programs, replacing human programmers.
- Management support, especially for planning technological change.
- Engineering support for design, manufacturing, and testing.

The space station not only provides an opportunity for using AI, but may well require its application for successful development. The knowledge base that will be created for designing, manufacturing, testing, and operating the space station needs to be planned from the start to accommodate AI techniques.

NASA should establish an AI research group at a NASA center close to an existing AI research community (Ames Research Center would probably be an ideal location). This group should have close ties with the AI research community, the areas of application within NASA, and NASA management.

To develop the AI systems that NASA will need in the years ahead, the AI research group could assist in retraining NASA system developers and work closely with NASA application-development groups.

This Page Intentionally Left Blank

II PREVIOUS NASA STUDIES

From time to time various individuals within NASA, as well as study groups that have included both NASA personnel and non-NASA people, have looked at the impact of future technology, including computer science and artificial intelligence. Because of the rapid development of AI as a technology, only the most recent studies are genuinely helpful. One landmark study, by the Outlook for Space Study Group (1976), recognized the potential impact of AI in "A Forecast of Space Technology: 1980-2000" (1976), but underestimated its rate of progress.

The NASA Study Group on Machine Intelligence and Robotics, chaired by Carl Sagan, met from June 1977 to December 1978. It concluded its deliberations by recommending a vigorous program of research combined with practical applications of machine intelligence and robotics in NASA.

In the summer of 1979, at Woods Hole, Massachusetts, the Innovation Study considered (among other topics) self-replicating systems. No official report was published, but the study's findings were subsequently discussed by Bekey and Naugle (1980).

Following a symposium on automation and future missions in space held at Pajaro Dunes, California, in June 1980, the 1980 NASA/ASEE Summer Study met at the University of Santa Clara, June-August 1980. They focused their attention on four space missions in the following categories: (1) terrestrial applications, (2) space exploration, (3) nonterrestrial utilization of materials, and (4) self-replicating systems.

The 1981 NASA/ASEE Summer Study met at the University of Maryland. Its principal theme was the utilization of computer science and technology in NASA, considered from the standpoint of NASA's needs as an organization, the present status of science and technology in NASA, and recommendations.

We shall make liberal use of the findings of previous studies. A bibliography at the end includes the published reports of such studies, as well as other references cited in the following text.

III OVERVIEW OF ARTIFICIAL INTELLIGENCE

Computers have been in existence now for almost forty years. While steadily decreasing in price, they have increased enormously in both usability and speed. During these decades of extraordinary technological development, computers have been used primarily for either "number crunching" (e.g., in numerical analysis) or file/data processing (e.g., in most commercial applications). However, because their inherent speed and accuracy enable extensive reasoning in domains that can be appropriately symbolized, computers can also be viewed as symbol-manipulating machines. For example, computer programs are available for symbolic mathematics (e.g., equation solving and indefinite integration) using algebraic manipulations of mathematical expressions rather than numerical techniques. The main advantage of such programs is that they can solve mathematical problems of great complexity far faster and more accurately than any human could hope to do. This ability to manipulate symbols opens up exciting possibilities for intelligent computer applications.

Artificial intelligence can be viewed as the study and development of techniques for exploiting this symbolic-reasoning potential. In practice, this means the building and testing of programs whose behavior can be called intelligent. This activity is combined with the goal of developing a science of cognition based on a symbol-manipulating, information-processing model of intelligence. While considerable progress has been made toward this general objective, most AI research concentrates on such specific aspects of intelligence as planning, reasoning, hypothesizing, learning, conjecturing, etc., as well as on the

engineering task of building systems and general tools. An understanding of how particular intelligent activities can be performed is expected to lead to an understanding of intelligence in general - its essential nature and processes.

The major problem in developing AI systems is in finding suitable symbolic representations of the domains of interest, along with commensurately appropriate inference systems to deduce any necessary conclusions. Expert systems are outstanding examples; these are programs that appear to the user as experts in their respective domains, allowing each user to interact with the system as if he were dealing with the expert himself (e.g., giving it information and asking it questions). Such expert systems use the same kind of reasoning as the expert, ranging from rough rules of thumb to precise deduction from the available data.

Today computer-based investigations of intelligence have diverged into two related subfields, one predominantly theoretical and the other predominantly practical. The first is cognitive science, concerned with the understanding and modeling of human thought. The second is AI, which is concerned with achieving computer-based intelligent behavior by any means. However, in practice, AI programs tend to emulate the way people think because, to ensure correct behavior, the researcher and user have to understand what the system is doing.

With the cost of computing power decreasing rapidly, the expense of programming the machine for a specific application has become the dominant overall system cost. Advances in AI have made possible great reductions in the cost of software development while increasing the flexibility and reliability of the resulting systems. In addition, AI systems greatly extend the range of useful activities that computer systems can perform, from handling data bases and number crunching to serving as interactive intelligent assistants and on-line intelligent experts. This enhanced capability is possible, in part, because recent rule-based AI systems make it possible to alter specific rules without upsetting the rest of the program. Also, these modular, rule-based representations are more easily understood by nonprogrammers.

Progress in artificial intelligence is limited more by conceptual difficulties than by hardware constraints. Generally, the hardest problem is the development of a suitable representation for each domain, and the accumulation of sufficient domain-specific knowledge (including search guidance) to enable useful inferences to be drawn in that domain. Despite these difficulties, AI has already produced many useful systems and further research can be expected to produce many more in the near future. These new systems will be significantly more powerful and adaptable than those currently available.

AI research combines general research into a broad spectrum of common problems with a more specialized kind of research focusing upon specific areas of application. The core of basic research is surrounded by various layers of applications (the "onion" model) as shown in Figure 1.

A. Basic AI Research

The core of basic AI research, as shown in Figure 1, consists of four components. These are described in the following subsections.

1. Search Theory

In problem solving, planning, theorem proving, game playing, and diagnosing, for example, an AI program must search through a large space of potential solutions. Usually this space is far too large to explore "by brute force," so knowledge - especially knowledge about the domain being investigated - is used to guide the search. This knowledge is often incomplete and is even occasionally in error, but the guidance it provides is considerably better than mere chance. Such inexact information is called heuristic. A considerable body of information now exists on how to conduct heuristically guided searches.

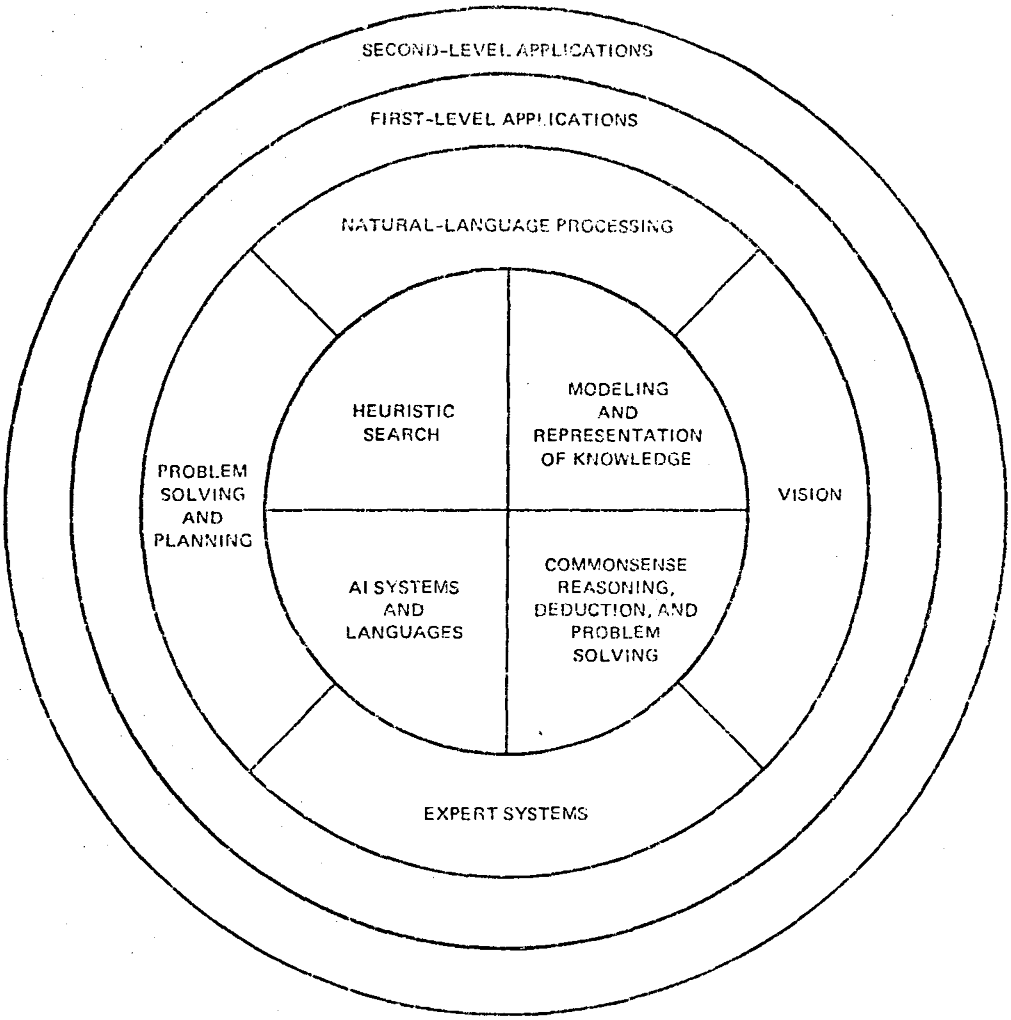

FIGURE 1 THE "ONION" MODEL OF ARTIFICIAL INTELLIGENCE RESEARCH AND DEVELOPMENT

2. Modeling and Representation of Knowledge

One of the most difficult tasks in developing an AI application is to devise a suitable representation and associated inference scheme. Such a scheme must capture the principal concepts of the domain and

allow the efficient computation of any required information. Sometimes representations can be studied in the abstract, but usually a representation and its inference procedures must be studied together in a given domain to ensure the efficiency of the resulting system.

3. Commonsense Reasoning and Problem Solving

AI research is developing methods for reasoning in uncertain and incompletely known situations. Such methods are necessary for understanding, planning and decision-making under real-world conditions. Success in this area is leading to systems that reason closer to the way people do, rather than working in the formal algorithmic fashion characteristic of most current programs, and so are much easier to build, understand, and modify. Because of their progress in developing such systems, practitioners are sometimes referred to as "knowledge engineers."

4. AI Programming Languages

Because of the different programming needs in building symbol-manipulating programs, compared with number crunching or data handling, many programming languages (and programming aids) have been developed for AI. In particular, LISP has been used by the AI community since its development in 1958. It has become the major AI language, with a powerful user environment that is unequalled in computer science. LISP, and packages written in it, have made possible the development of very complex AI systems. LISP has also found many uses outside the bounds of AI. Recently new AI programming languages, such as PROLOG, have become available. These new languages offer greater programming convenience than LISP for certain applications.

B. AI Applications

The next layer of AI research surrounding the core is concerned with areas that benefit from progress in basic AI research. Many AI laboratories around the world have groups working on the following major applications.

1. Expert Systems

This is the name given to a class of AI systems that utilize a significant amount of expert information about a particular domain to solve problems in that domain. The feature distinguishing such expert systems from standard application programs is that the expert information is stored in a transparent form (usually rules) that is used by the system in accordance with and in response to the current problem. All current expert systems separate the expert information from the inference procedure that reasons on the basis of this information. Typical numerical-application programs (e.g., for space navigation) dictate in advance how a computation is to be performed for every possible computational step, and thus require considerable modification if the system needs to be changed. Another difference between expert systems and conventional computer science is that expert systems are designed to be able to reason on the basis of incomplete and uncertain knowledge.

Expert planning systems are still in a state of development. These systems can produce complex plans that obey many constraints with a potential for achieving an accuracy and speed that humans are not able to match. Although building such planning systems entails a high overhead cost because of the large amount of knowledge that must be acquired and the need to debug (verify) this knowledge, the potential return is substantial because such systems can be used repeatedly at many different locations, and without additional cost. DEVISER, developed at JPL, is such a planning system; it can be used for such tasks as planning missions. With the addition of suitable domain knowledge it could be used for designing vehicles and planning the control of spacecraft.

2. Vision

The amount of information generated by an imaging sensor (e.g., a TV camera) is enormous, so that the task of reducing this information to a machine-understandable internal representation raises many challenging problems. Vision systems that recognize limited classes of images are already commercially available, but considerable work remains to be done before AI vision systems can match the capability of the human eye. A large body of information exists regarding methods for extracting relevant features from an image (e.g., edges, regions, shadows, etc.), but the higher-level problem of integrating this information for the purpose of recognizing objects has been solved in only very limited circumstances. The low-level vision processing will require high-speed, special-purpose VLSI chips to perform feature extraction in real time.

3. Understanding Natural Language and Speech

Here the problem is to reduce the input to an internal representation that captures the intended meaning. In speech understanding the input is the digitized acoustic signal, while in natural-language understanding the input is text usually entered from a keyboard. While many problems remain to be solved, useful natural- language systems are already commercially available. However, the problem of understanding speech is much more difficult because of signal interpretation difficulties. Despite these obstacles, an experimental speech-understanding systems with a vocabulary in excess of a thousand words has been developed.

4. Automatic Programming

The goal in automatic programming is to transform a given program specification into a working, provably correct program for a particular computer. When the desired program is only informally stated (e.g., by sample inputs followed by their corresponding outputs), then the task of automatic programming constitutes inductive learning; there is no assurance that it will produce the intended program. A formal program specification can be transformed into a working program more easily. Programs of moderate complexity have already been created.

5. Problem Solving and Game Playing

Considerable progress has been made in these areas, which have been subjects of study since the earliest days of AI. For example, chess-playing programs now regularly defeat all but grand masters, and problems such as Rubik's cube are routinely solved by problem-solving programs. Such idealized domains were studied because they explicate difficult fundamental issues in AI research. These domains enable research progress to be made on the interesting fundamental issues, unencumbered by the distracting complexities of the real world. Game playing is a major analogical context for understanding human behavior and planning in multiagent situations.

6. Robotics

For a robot to achieve goals and cope with unexpected conditions, it must build and maintain a world model. This model, used to evaluate the effects of proposed actions (plans), depicts the expected state of the world so that sensors can detect error conditions. As a robot interacts with the world, sensors can be used to update the world model and guide future actions.

7. Learning

This general term covers any AI research involving a program that improves its behavior as a result of experience. Many interesting learning programs have been created, but much research remains to be done before learning can be regarded as a routine part of an AI system.

C. Applied AI

The outer layers of AI applications in the onion model include systems of increasing specificity for use in particular areas. These systems are characterized by a concern for reliable performance and ease of use, unlike AI research programs that tend to be usable only by their originators. Areas in which such AI systems have been developed include medicine, geology, automatic control, signal understanding, chemistry, and others. A major limitation upon the growth of applied AI systems is the shortage of trained AI personnel (knowledge engineers) to work with experts and transfer their knowledge to a useful working system. The availability of high-level languages for construction of rule-based reasoning systems, such as OPS5 and PROLOG, will make this task easier, as recent experience at DEC (McDermott, 1982) has shown.

Recently, much AI research has been directed at expert systems, as most AI applications involve some expertise. These systems usually employ rule-based reasoning because of the advantages of representing knowledge in this form. For example, when an expert system applies the relevant rules to a particular situation, a record of such rules and their context can be kept. This record can be used by the system to explain its reasoning to the user - to whatever level of detail may be desired. In addition, if the expert building and debugging a system does not agree with the system's reasoning in a specific situation, he can usually resolve the discrepancy locally by modifying an existing rule or adding a new one. Since the rules are largely modular, such modifications normally do not alter previously correct inferences.

A major concern of expert-system research is how to use uncertain rules (rules of thumb), so that degrees of belief in individual rules can be combined to obtain a final degree of belief. Since most human reasoning takes place on the basis of uncertain information and weakly supported implications, it would clearly be of great advantage if expert systems could use similar types of inferences. Rather than trying to imitate human behavior directly, one line of current research is directed toward formal models for reasoning under conditions of uncertainty that utilize all the information available, but without overstating the strength of any single conclusion. Practical systems that can make probabilistic judgements given particular information, are available and are expected to improve rapidly.

Although artificial intelligence has made great advances and has already produced a number of practical systems, future AI systems can be expected to grow significantly in scope and capability. This growth will be especially expedited by the emergence of special-purpose AI machines (e.g., PROLOG and LISP machines being developed as part of the Japanese fifth-generation project). However, such machines will not themselves create powerful AI systems, as the main limitations are still conceptual. Considerable basic research remains to be done, particularly in the area of knowledge representation and associated inference procedures. As some of this research is especially pertinent to NASA's requirements, it should obviously be accorded commensurate attention and emphasis by NASA.

IV OVERVIEW OF NASA APPLICATIONS OF AI

Some previous studies have looked at possible applications of artificial intelligence in NASA - most notably, the NASA Study Group. That group found prospective applications in smart sensing, manipulators, control systems, spacecraft crisis analysis, locomotion systems, assembly of large structures, vision systems, software tools, and computer networks. The number of potential applications appears to be very large indeed.

The impact of AI on NASA will be similar to the impact of computers over the last twenty years, except that AI will probably develop faster. In the case of computers, the first use was in a few high-priority, operational applications in which computers were not only obviously advantageous, but for some purposes even essential. Computer resources were limited and concentrated, as were the personnel who used them. However, computers have become ubiquitous and personnel skilled in their use are found throughout the NASA organization.

AI is just now beginning to be used in business and industry, with the first truly practical applications implemented during the past year or so. The natural-language interface INTELLECT has been sold and installed in over a hundred locations. The R1 system, developed by CMU and used by DEC for configuring VAX computers, is probably the first expert system to have been put into industrial service. Industrial robots in a few large manufacturing organizations are being equipped with devices such as vision modules (from the AI research laboratories) to make them "smart." Expansion and proliferation are already on the

horizon, however, as many organizations are now engaged in developing systems either for their own use or to sell to future users. Even if NASA itself did nothing about R&D in artificial intelligence, it would inevitably become a major user of AI products along with business and industry.

The Space Study Group, in its report to the NASA Administrator, "Outlook for Space" (1976), divided technologies among three categories, depending on the manner of their support:

- By industry, with little or no government support.
- Partially by industry, with a reasonable fraction of NASA support.
- NASA-supported because of unique pertinence to space requirements.

The Space Study Group put AI in the second category, stating that "There could be major advances in automated (machine) intelligence, enabling spacecraft and surface rovers to conduct important tasks or sequences of operations under human direction, but without the need for constant step-by-step human control."

Similar study groups formed since the Space Study Group's report have concluded that NASA will probably find AI to be an absolute requirement in carrying out future missions, such as the space station currently in the planning stage. In general, AI will be used in two kinds of applications: (1) to improve existing functions, and (2) to enable the performance of tasks that were formerly considered impossible.

As regards the improvement of existing functions, AI can have its greatest impact in replacing human workers engaged in the handling and processing of information. For example, a good prospect for early application of an AI-type planning system is the task of improving the productivity of personnel in mission planning and related activities. One system, DEVISER (a JPL development), is being considered for use in several NASA centers.

As regards the enabling of heretofore impossible tasks, AI will be required to support advanced aeronautical and space projects. For example, AI techniques for fault isolation, diagnosis, and repair will be needed to achieve the degree of reliability required in future aircraft and spacecraft.

Examples of NASA applications can be found for each of the major application areas of AI that were described briefly in Section II:

<u>Expert Systems</u> - Expert planning systems will apparently be the first AI systems used in NASA, since one of them, DEVISER, is already in an advanced stage of development. It has the potential to be used at several places in NASA and to provide a foundation for future systems with even broader scope.

<u>Vision</u> - AI techniques will be a vitally important aid in processing the enormous amount of imagery anticipated from planned missions, such as LANDSAT.

<u>Natural-Language and Speech Understanding</u> - Probably the most available AI technique today. Immediate applications in NASA are possible as natural-language interfaces to data bases.

<u>Automatic Programming</u> - Tremendous long-range potential for replacing human programmers. AI-type programming aids can increase the productivity of programmers in the near term.

<u>Problem Solving and Game Playing</u> - A useful research tool for devising and developing new AI capabilities.

<u>Robotics</u> - In the form of teleoperation, already recognized as an important NASA technology. As missions become more complex and numerous, more autonomous robots will be needed in response to rising costs.

In summary, the number of potential applications of AI in NASA is very large and pervades all of NASA's activities. The use of expert planning systems such as DEVISER in mission planning would appear to be a reasonable choice for an initial application. As experience is gained, the scope of future applications can be increased. Because of

the inevitable widespread use of AI in NASA, its initial applications will not only provide additional detailed technical knowledge and an increased reservoir of skilled workers, but will also establish the management framework needed to guide the further development of AI systems in NASA.

V APPLICATIONS OF AI

A. The Scope of AI Applications in NASA

Although NASA has begun development of a few important applications, widespread utilization of AI by NASA is expected in the near future. What we are now witnessing is just the inceptive phase of a new technology that is just starting to emerge from the research laboratory. Any discussion regarding eventual applications involves the kinds of uncertainties usually encountered when one ventures to predict the future. Some of the applications discussed here may turn out to be more difficult than anticipated, whereas others, including some ultimate successes, may be overlooked entirely. Fortunately, several previous studies that have already considered some possible applications of AI in NASA can serve as a foundation for further conjectures and proposals. However, the reader should remember that an attempt is being made to suggest possibilities over the entire broad spectrum of NASA activities. Inevitably certain specific future applications, including some that might later prove to be important, may not be mentioned at this time.

Just as computers have proliferated during the past two decades, so will the applications of artificial intelligence. This will be occurring at the same time that other technologies, such as very-large-scale integrated semiconductor circuits and digital computer networks, are being deployed. Just as robotics, a part of AI technology, is changing factories, so AI will change offices. The pervasiveness of potential AI applications will force organizations such as NASA to make strategic choices in the next decade promoting the use of AI technology.

To put possible applications in perspective, it would be helpful to consider NASA's missions (as listed in the amended version of the National Aeronautics and Space Act of 1958). These are as follows:

- o Atmospheric and space science
- o Aeronautical and space vehicles
- o Leadership in aeronautical and space science and technology
- o Technology transfer
- o International cooperation
- o Ground propulsion systems
- o Automotive propulsion systems
- o Bioengineering

All of NASA's activities support one or more of these missions. A direct relationship is evident between AI and the second mission, aeronautical and space vehicles, e.g., in autonomous spacecraft and as stressed in such previous studies as the one conducted by the NASA Study Group. In addition, an examination of each of the other missions would reveal important prospective applications for AI. The use of AI techniques for understanding imagery from a geostationary operational environmental satellite (GOES), for example, could typify AI's practical value in atmospheric science.

Future aeronautical and space science and technology are expected to become increasingly dependent on artificial intelligence - to such an extent that technological leadership in aeronautics and space science will presuppose leadership in AI. To succeed in accomplishing the third mission, therefore, NASA must lead in AI, at least in those aspects of it that are most relevant to NASA's requirements.

The fourth mission, technology transfer, would benefit national security, other government functions, and the private sector by spinning off from NASA the AI technology developed for NASA applications. As

noted later, AI itself, in the form of expert systems, could facilitate the process of technology transfer.

B. Information Management

The acquisition, storage, transmission, and dissemination of information pervade all of NASA's systems and programs. These vital functions have received the attention of several studies in recent years.

1. NASA Study Group

The NASA Study Group looked at a large number of potential applications of AI in NASA, including all the above mentioned aspects of information management. The final report of the Study Group listed the following areas of technological opportunity:

a. Robotics - Use of AI in autonomous exploratory spacecraft and for an autonomous lunar explorer.

b. Smart Sensors - Use of AI in image understanding for such applications as remote sensing, crop surveys, and cartography. Also envisaged is the use of AI vision in teleoperators and a Mars rover.

c. Mission Operations - AI in mission monitoring, and in sequencing and control.

d. Software Development - Use of on-line, interactive facilities for software development; other aids to programmer productivity.

e. Man-Machine Systems - Low-level, detailed control by intelligent computers, leaving humans free to carry out higher-level, more complex decision-making and administrative responsibilities.

Not all of these applications mentioned by the NASA Study Group are AI, strictly speaking. Software development, for example, is descriptive of the technical environment in which AI research is conducted. Nevertheless, those applications of AI that are identified are significant.

2. 1980 Summer Study

The 1980 Summer Study considered a hypothetical earth-sensing information system that could be implemented during the next two or three decades and that would require intensive application of artificial intelligence techniques. This system was studied because of its potential economic benefits, and because present earth-sensing systems are collecting data at a rate far exceeding the rate at which the information can be digested (formatted, cataloged, verified, annotated, calibrated, etc.) and disseminated. Called the Intelligent Earth-Sensing Information System (IESIS), its demands would be at the far extreme of present NEEDS (NASA End-to-End Data System) activities, particularly in the planning, scheduling, and control of satellite systems.

In the report of the 1980 Summer Study, the IESIS is described as follows:

> An intelligent satellite system that gathers data in a goal-directed manner, based on specific requests for observation and on prior knowledge contained in a detailed self-correcting world model. The world model, or knowledge base, would be a part of the AI system, eliminating the processing and storage of redundant information.
>
> A user-oriented interface that permits natural-language requests to be satisfied without human intervention from information retrieved from the system library or from observations made by a member satellite within the system.
>
> A medium-level onboard decision making capability which optimizes sensor utilization without compromising user's requests.
>
> A library of stored information that provides a complete detailed set of all significant Earth features and resources, adjustable for seasonal and other identifiable variations. These features and their characteristics are accessible through a comprehensive cross-referencing scheme.

Typical data-processing steps in a user-IESIS interaction are listed in Figure 2.

DATA ON DEMAND

USER
LOGON
SYSTEM/USER INTERFACE
–NATURAL-LANGUAGE QUERY-RESPONSE
ON-GROUND PROCESSING
–KNOWS ALL SATELLITES
–KNOWS ALL USER REQUESTS
–KNOWS ALL PRIORITIES
–BUILDS PLAN FOR IMPLEMENTATION
–GENERATES COST ESTIMATE PRIOR TO PLAN EXECUTION
UPLINK
SATELLITE SENSING AND PROCESSING
–RAW DATA COLLECTION
–MAKES DECISION ON DATA COLLECTION
–ONBOARD PROCESSING
DOWNLINK
ON-GROUND PROCESSING
–PICTURE PROCESSING
–UPDATE WORLD MODEL
–ARCHIVE
–DELIVER RESPONSE TO USER
USER

FIGURE 2 DATA-PROCESSING STEPS IN A USER-IESIS INTERACTION

The 1980 Summer Study did not describe all the applications of AI that would be required for the IESIS, but noted a few of the important ones.

Natural Language – Primarily to provide novice users with direct access to IESIS. The natural-language interface for IESIS requires a domain model, a user model, a dialogue model, reasonable default assumptions, common world knowledge, and explanatory capability.

Problem Solving – Two specific issues were considered. The first of these was the communication link capacity. An AI planner could significantly reduce all communication link volumes by taking several factors into account: each individual user request, the ensemble of all current requests, and a forecast of expected requests. The second issue considered was the number of satellites. If the IESIS concepts for observation can be implemented, it may be possible to shift most data-acquisition tasks to a lower level of resolution.

This system would allow data-acquisition by orbiters at higher altitudes with greater fields of view: thus, a smaller number of satellites be required. Success would depend upon a high-level capability to understand relationships between sensor readings and the actual state of the world as by human-oriented descriptors. This is precisely the problem in visual perception that is addressed by AI.

The foregoing observations not only propose a feasible long-range goal for an earth-sensing information system, but, at the same time, point to some of the key AI techniques that will be required in such a system. One can assume, moreover, that other NASA information systems would need some of the same AI techniques - natural language, for example - for similar purposes. Space exploration, certainly, presents many problems of control and orientation that are prime prospects for AI applications.

3. 1981 Summer Study

The 1981 Summer Study considered techniques for knowledge acquisition and dissemination, for storing, retrieving, and processing complex and/or large volumes of data. The 1981 Summer Study report concluded that assimilation of the vast quantities of information to be collected by NASA will be impossible without new techniques for massive data storage, data-base management, and data processing. The report states that NASA will need sophistication in knowledge representation well beyond the current state of the art.

Applications of expert systems are discussed in an appendix to the report. It includes the comment that current procedures for tracking, command sequence generation, and data processing are labor-intensive, and that their automation on a wide scale seems a practical goal with current AI technology. The effort now underway at JPL to automate the process of uplinking commands to a spacecraft is an important first step. Expert systems can increase the efficiency of command sequence generation and verification, as well as the overall control of uplink activity management. The appendix also notes that knowledge-based

planning systems can be used to improve the scheduling, monitoring, and control of operations such as those involved in the shuttle turnaround.

The report noted the following potential applications of expert systems in information acquisition, processing, and dissemination:

- Ground support operations for the space shuttle and for earth orbit and planetary spacecraft.
- Shuttle scheduling and monitoring for the shuttle turnaround.
- Fault analysis and management for the shuttle and other space systems.
- Monitoring and management of astronauts' health.
- Scheduling of meetings.
- Software development aids: tools for programming and for software project management.
- Aids for research and analysis, e.g., DENDRAL and MACSYMA; experiment-scheduling, design, and interpretation aids; sophisticated real-time control of experiments.
- Image interpretation; anomaly detection for further analysis by humans.
- Automatic management of technical documentation.
- Self-managing data bases, i.e., data bases that can reorganize themselves as needed for efficient access, check the validity of inputs automatically, etc.

Thus, the three study groups (NASA Study Group, 1980 Summer Study, and 1981 Summer Study), each in a different way, have looked at the need and possibilities for AI applications in the acquisition, storage, transmission, and dissemination of information.

4. Mission Planning

As already noted, JPL has begun the development of an AI planning system (DEVISER) for generating command sequences for the uplink, an

application identified by all three groups. It appears to be a sensible choice that, at the same time, emphasizes an important point. The first application of AI in NASA is a planning task, whereas those that have been first outside NASA have been either natural-language or expert systems. But NASA, of course, has its own unique priorities that cannot be equated with those that obtain elsewhere.

5. Other Applications

The first applications of AI in NASA, as exemplified by DEVISER, are in the uplink process. An earlier study (Brown, April 1981) considered using AI in Mission Operations at Goddard Space Flight Center. The following applications were recommended:

- An expert system to aid the spacecraft analyst, augmented by a natural-language interface and a planning system.
- A planning system for mission planning in the Payload Operations Control Centers (POCC).
- Use of AI techniques in the Command Management System (CMS).
- A planning system for scheduling of resources in the Multisatellite Operations Control Centers (MSOCC).
- A planning system for scheduling in the Network Control Center (NCC).
- A planning system for scheduling jobs n the Image Processing Facility (IPF).
- AI techniques for image understanding in the quality assurance function of the IPF.
- Use of AI techniques for image understanding in the IPF's automatic registration of ground control points.
- Use of AI techniques for image understanding in satellite imagery, e.g., in crop surveys.

Other possible applications of AI are being considered at various NASA centers. For example, AI is being investigated at Johnson Space Center, Flight Control Division, as a means of monitoring the space

shuttle downlink for irregularities. They are also looking at the possible use of AI in connection with a DBMS for a system of distributed data bases. Under a contract with Goddard Space Flight Center, the MITRE-developed, knowledge-based system KNOBS is being adapted to demonstrate its ability for doing crew activity planning for the space shuttle, with a view to its possible use with the space station. The KNOBS adaptation, called KNEECAP (Knowledge-Based English-Enquiry Crew Activity Planner), MITRE Project 8980, is being developed in cooperation with the Operations Division at Johnson Space Center.

A planning system similar to DEVISER could be used for scheduling science experiments on the space shuttle. This is now done manually at Goddard Space Flight Center, in the Engineering Directorate.

The huge amount of data generated by imaging systems creates a massive information storage problem. Artificial intelligence makes it possible to control this output by examining and comprehending it in real time, then selecting and transmitting only that information considered important according to certain specified criteria. This would enable monitoring satellites to detect and transmit interesting events, ignoring all others, and to direct onboard sensors to observe important events as they take place. Because of the transmitting delays inherent in deep-space missions, such understanding/adaptive systems must be autonomous.

The large volume of data and the critical need for accuracy and error-free processing make intelligent expert systems extremely desirable for planning scientific experiments and processing the data. When such expert systems are introduced, they not only aid the development of those branches of science in which they are employed (e.g., MOLGEN) but actually engender a new kind of communication and cooperation within the community of participating scientists.

AI expert systems could aid the process of technology transfer by providing on-line expertise in a particular area, e.g., in fuel cells.

In this case, the expert system would include a knowledge base containing the knowledge of human experts about fuel-cell technology and its applications. Users would then obtain access to the expert system via a computer network, which is similar to the way librarians now access the NASA RECON system for bibliographic information about NASA technology. In the case of an expert system, however, the user would effect technology transfer by interacting directly with the expert system in the language of his discipline.

6. Summary

The area of acquisition, storage, transmission, and dissemination of information offers a rich harvest of possibilities for applications of artificial intelligence. The first applications are in the uplink process, an apparently reasonable choice, even though NASA's long-range plans for AI are still being formulated. A long-range goal and a well-defined plan will be essential for managing the development of AI applications. Otherwise the alternative may be automation of a number of discrete little fragments of the present system that may not constitute optimum or reasonable steps toward the desired goal or even be compatible with one another. The initial approach suggested above - starting with an application like DEVISER for planning command sequences and then expanding the scope of the system - appears to be both reasonable and timely.

C. Engineering

NASA has substantial investments in engineering systems for the design of aircraft, spacecraft, etc. Some applications of AI that are being developed in research laboratories are expected to influence and become a part of such systems.

1. 1981 Summer Study

The 1981 Summer Study assessed NASA's ability to design, model, develop, test, operate, and manage large, complex (possibly autonomous) systems. Two facets of engineering systems were recognized. The first of these concerns the technologies needed for a particular kind of system. AI, for example, is an essential technology for building autonomous spacecraft. The second facet comprises the collection of methods, tools, management techniques, and design procedures employed in engineering systems. Here too, AI techniques are viewed as a requirement.

The same study group also advocated a method of computer-aided systems engineering (CASE) in which the central component is a complete description of the system in machine-readable form. The life cycle of a CASE system is shown in Figure 3. CASE would have a natural-language interface, similar to English, that would make use of AI. The machine-readable description of the system would include all the reports, working papers, correspondence, studies, models, data sheets, analyses, specifications, parts lists, drawings, plans, schedules, costs, and supporting information that pertain to the system. This knowledge base would also interface with and support the use of computer aids for all phases of development, testing, and operation, such as computer-aided design (CAD) and computer-aided manufacturing (CAM).

The representation of the system's knowledge in a useful form needs to be based on AI techniques for knowledge representation. Then AI techniques for design, diagnosis, and other engineering activities can work in conjunction with the same knowledge base.

The CASE methodology proposed by the 1981 Summer Study has a number of potential advantages. Automation of engineering documentation is expected to reduce the time spent in labor-intensive activities and lower engineering costs. Other features are described in the 1981 Summer Study report as follows:

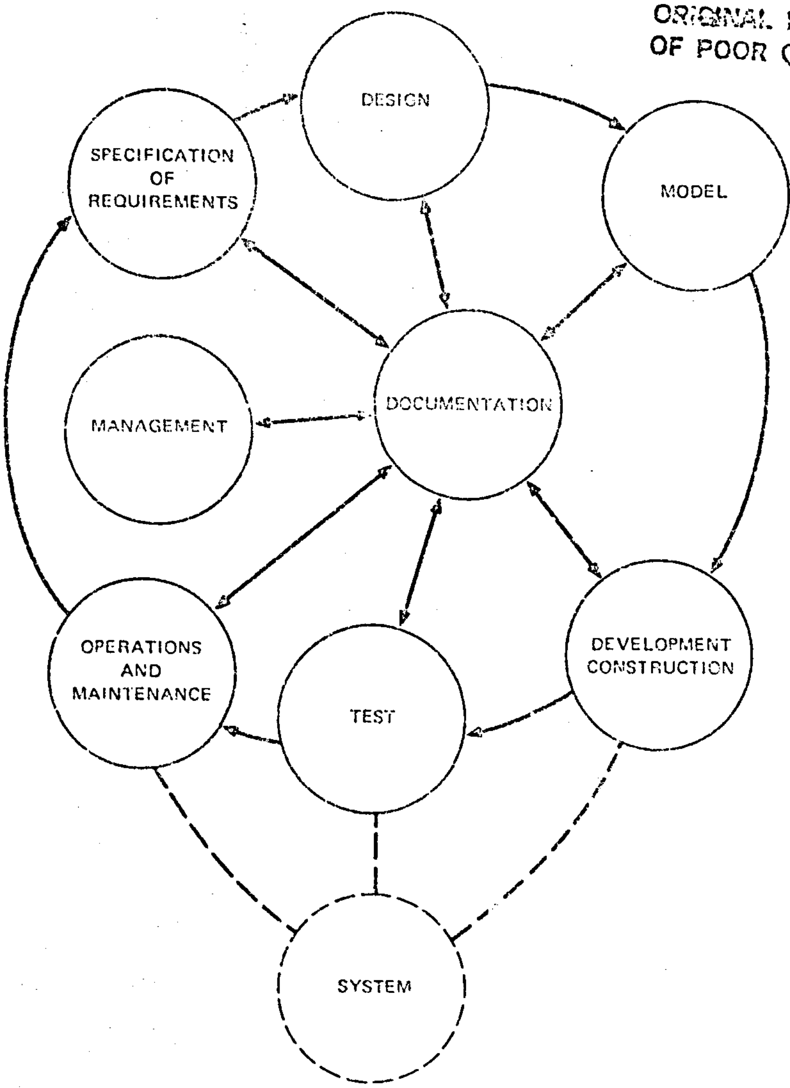

FIGURE 3 COMPUTER-AIDED SYSTEMS ENGINEERING

CASE has implications for personnel mix and talent pools in the Agency. A new breed of technician, as e.g., "Integration engineers" will be needed to provide consistency and constraint checks, since all requirements must be met by the design procedure. Components which appear to satisfy no requirement must be scrutinized. Properties and limits of materials have to be checked against how they will be used. Interfacing of components within a system could be checked by the demon engineers for compatibility. Modeling and simulation of systems

will become standard and will be used to reveal incongruities with all modifications undergoing extensive checks to determine the impact on other parts of the system. Cost estimates could be automatically generated at any point during system development. Under these conditions, designs are never considered frozen and thus yield more flexible and reliable systems.

After an original unit under development becomes operational, CASE design products should remain on-line and available for the life of the particular integration. Higher-level CASE functions will permit designers to monitor system-wide development, testing, and operations. For instance, CASE can gather historical statistics, identify procedural or developmental bottlenecks, and pinpoint specific trouble ares where redesign could pay off in improved performance. CASE could also be used to diagnose faults and repair deficiencies by formulating and modeling a fault hypothesis and comparing results with the real situation.

Finally, CASE vastly simplifies and enhances system management. Documentation is central to the process, so management functions such as scheduling and monitoring can more easily be conducted by this means. Queries about the system should provide exactly the information requested by the user. CASE should even be able to refer the user to the technical literature for more extensive explanations of the theory of the particular parts of the system. Such features will allow managers to continuously monitor and guide the development and operation of new systems in real time.

2. MIT ARAMIS Study

In a contract study for Marshall Spac Flight Center, MIT recommends a similar approach (Miller et al., 1982).

NASA should consider developing a computer simulation and data management system for satellites, to be implemented end-to-end, i.e. from the original mission definition, through spacecraft design, manufacture, test, integration, launch, on-orbit checkout, nominal operations, spacecraft modifications, and fault diagnosis and handling. Such a system would enhance communication between mission supervisors, and reduce documentation costs. As the study group found in its own data management system, important objectives are that each individual user should have access to all the data, and that paper should become secondary to the computer as a communication medium.

3. IPAD

An effort along these lines has been under way for a number of years, sponsored by Langley Research Center and known as the Integrated Programs for Aerospace-Vehicle Design (IPAD) (Fulton, 1982). This effort is coordinated with the Air Force's Integrated Computer-Aided Manufacturing (ICAM) program; the two programs are being developed jointly, with some common elements. Both can be regarded as directed toward and, in effect, validating the CASE methodology described by the 1981 NASA/ASEE Summer Study. Figure 4 illustrates the ICAM program.

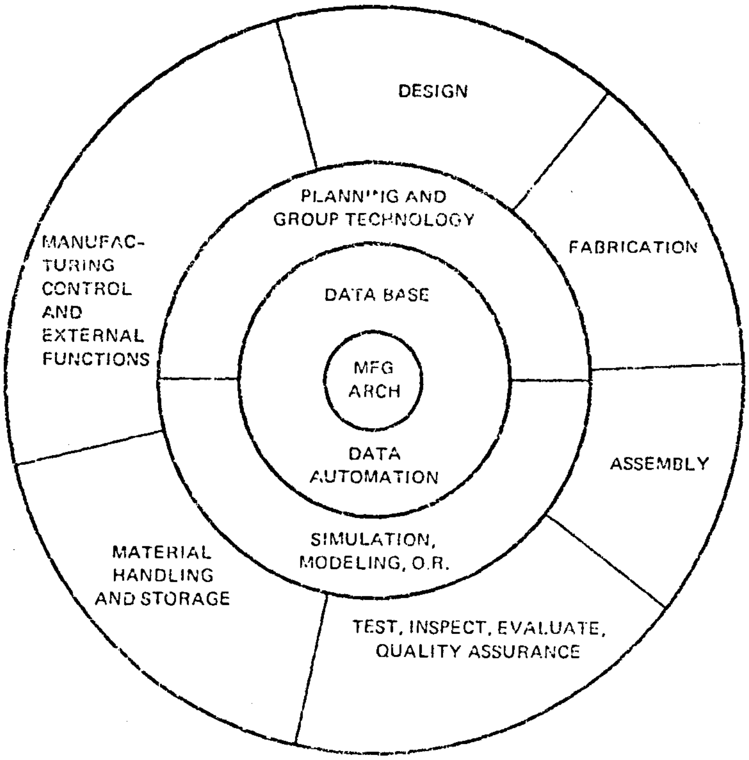

FIGURE 4 ICAM LOGO

The joint industry-government IPAD program is a general-purpose interactive computing system to support engineering design, with significant capability for managing and manipulating engineering data. It is intended to support activities at all levels of design - conceptual,

preliminary and final - and to aid in the assembly and organization of design data for manufacturing. It is envisioned as being implemented on a high-speed network of computers from various companies. IPAD efforts have been concentrated on data management for design and manufacturing. In the future, AI-type tools can be developed to work with the IPAD knowledge base for design, diagnosis, scheduling, and other tasks, according to Robert E. Fulton, manager of the IPAD Project at Langley Research Center.

4. AI in Future CAD/CAM

Although CASE is still hypothetical and neither IPAD nor ICAM is yet fully developed, they share the same basic concept and approach to systems engineering. The concept and some of the elements of the IPAD and ICAM programs that have been developed can be used by NASA for such systems as the space station. Strong central management will be required to enforce the standards that will be required for this approach, but the benefits that can be realized later when more elements of IPAD and CASE become available will be important for the success of future missions.

CAD and CAM are widely used in NASA and in those sectors of industry that work with NASA, particularly the aerospace industry. JPL and most of the NASA centers have substantial investments in CAD/CAM hardware/software, including computer aids for the following areas and purposes: structural analysis (NASTRAN), modeling for analysis of lattice structures, analysis of test data, electrical layout, computational fluid dynamics, aircraft dynamics, modeling in theoretical physics, spacecraft configuration modeling, design of printed-circuit boards, design of infrared and advanced sensors, LSI/VLSI design tools, structural analysis of large antennas, analysis of vehicle geometry (aerodynamics, performance, mass/size, and cost), vibration analysis, structural and mechanical design of propulsion systems, scheduling, assembly, and parts inventory.

CAD/CAM usage is expected to grow apace while, at the same time, becoming more integrated as computers and computer networks proliferate. As a concomitant of this process, AI techniques will be incorporated into what is now called CAD/CAM. The basic AI approach to the representation of knowledge will influence the design of the data bases and their management. AI techniques for natural-language interfacing, vision, planning, and design will become part of the computer aids available to engineers. The increasing complexity of NASA systems, straining or exceeding the limits of human comprehension, requires new and better methods to manage the complexity. The capability of artificial intelligence, both existing and potentially realizable, offer effective solutions to this problem.

5. Space Station

An outstanding opportunity, perhaps a necessity, exists for the application of computer-aided engineering to the space station. The knowledge base that is now beginning to grow with the Space Station Task Force and the concept studies being conducted by the aerospace industry needs to become part of the space-station knowledge base that will be used throughout its lifetime, from original concept to perpetual maintenance. As such, it must be guided not only by the essential standards for compatibility, such as standard character codes, but also by the principles of knowledge base construction that are being discovered, tested, and formulated in the course of research in artificial intelligence. Many data bases that have been or are being developed for non-AI applications can be incorporated in AI systems without difficulty (parts lists, for example). However, when knowledge (e.g., decision rules) is the basis for decisions that may require explanation or modification, such knowledge must be represented in a way that will facilitate the needed explanation or modification. If this is done for data bases connected with the space station, a good foundation will then have been created for other space station subsystems that will utilize the original space-station knowledge base and embody artificial intelligence.

D. Institutional Management

1. Applications in Traditional Management

Many opportunities to apply AI profitably will appear in management, probably in all the traditional management areas: goal-setting, decision-making, policy formulation, evaluation, planning, budgeting, accounting, auditing, personnel management, training, career development, legal affairs, litigation, investigations, contracting, procurement, interorganizational relations, intelligence, etc. Just as computers, word processing, and office automation are now affecting all these functions, so also will AI have an impact.

Applications can be foreseen at any organizational level and for any activity in which access to information and its organization and analysis are essential elements. Natural language will be particularly useful in situations in which operations are intermittent and users inexperienced. Expert systems meet an entirely different set of needs by offering greater productivity in several ways:

- Replacing human intelligence in the performance of relatively simple repetitive and time-consuming tasks, e.g., automating the desk worker who routinely processes papers.
- Extending the availability of specialized knowledge, with fewer delays and at less cost than having to rely on scarce human experts.
- Applying knowledge-based control, organization, and strategy to human-machine interaction.
- Verifying or supporting complex analyses and decisions.
- Estimating costs and making budget projections.

In a broad sense, these systems may be considered "people amplifiers" that do not replace humans, but rather augment them and expand their capabilities.

2. The Changing Organization

The 1981 Summer Study discussed the impact of computer science and technology on NASA as an organization, noting that as the state of the art of computing changes, so will the nature of the organization. As a result of technological developments achieved many years ago, computers have already affected most organizations, including NASA. Technological progress is hardly slowing down - in fact, it is speeding up. Advances in semiconductor technology, for example, are accelerating the changes and proliferation of computing, and will continue to do so for some years in the future. Along with this whirlwind of technological change, or, to put it more accurately, as an integral part of it, AI is beginning to move from the research laboratory to its first applications in the real world.

a. The Role of AI

The 1981 Summer Study also noted that the same technology that has such an impact on the organization can also be used to manage its introduction. Thus, one area of potential AI application is in the management of change itself. Technological change is the primary concern, but the techniques are applicable to changes regardless of their nature or cause.

Briefly, the 1981 Summer Study outlined a method of adaptation to change that included the following:

(1) Knowledge acquisition throughout the organization and simultaneous development of an organizational knowledge base.

(2) Planning for organizational change.

(3) Plan execution.

(4) Evaluation.

(5) Replanning - then return to (1).

Most of these activities can be facilitated or performed more effectively by using AI systems.

b. Organizational Change

Because technology changes the way jobs get done, the organization must change accordingly. Unless it succeeds in doing this, it cannot expect to realize the full benefits of the new technology - in the worst case, it may not be able to adopt the technology at all. The success of NASA's missions depends not only on its readiness to use new technology, but on its ability to make whatever organizational changes may be necessary to adapt to such technology.

c. AI-Type Planning Systems

Although AI-type planning systems have not yet reached the same stage of development as expert systems, they show promise of becoming a very useful kind of AI application, and are beginning to receive more attention in the AI community. In this case, managing NASA's adaptation to change, an important advantage of the AI planner over its human counterpart is its ability to handle a much larger and more complex knowledge base. Such an AI planner would be used at many different levels, but because the AI planner can handle hierarchies of plans with their attendant details, the structure of the organization would also change. That is, fewer persons would be required to manage an activity, since the human manager, augmented by the AI-type planner, would be able to handle more details, together with their complex interactions.

Planning takes place at many levels and for many different purposes. Planning for change is often considered to be a top-management job, a job that tends to be neglected at all levels of management - at least in traditional organizations that have been stable or relatively crisis-free. Indeed, studies of organizations show that those that are

stable and successful resist change, whatever the outside motivating force. They exhibit a kind of homeostasis. Only in a perceived crisis, when survival is at stake, do organizations accommodate change willingly. This can be fatal in a rapidly changing world where, by the time the crisis is perceived, it can be too late to adapt (e.g., the U.S. automotive industry). (This is a generalization that admits of some exceptions.) However, awareness of the need for change can lead to planning for it and, eventually, to change itself. Ideally, the awareness should come from the top of the organization, where an overall strategy can be developed. That strategy (another word for plan) must be based on the considerations already stated: awareness of the need for change, assessment of the prevailing situation and of change-impelling forces (knowledge acquisition), determination of future goals or directions, and planning. As already noted, AI is concerned with the functions of knowledge acquisition and planning.

A strategy promulgated at the top of the organization should decide on the policies that will govern the development of AI applications for organizational change. Some key questions will have to be asked: What organizational planning activities have priority? At what levels should AI first be applied in these applications? How much funding should be committed to these AI planning applications? Answers to these and other questions will depend both on the NASA missions and on what AI is realistically able to accomplish. AI planning systems will be used by NASA first for planning missions, not for the planning of change. JPL's planning system, DEVISER, is being developed for planning Voyager missions - specifically, for generating the commands to be sent to the spacecraft. Experience with DEVISER and other planners to follow it will provide the foundation for applications of AI planners at higher levels of management - for example, in planning for technological upgrading of the space station.

Although the strategy governing applications of AI planners for change in NASA must come from the top, implementation may take place at any level. At whatever level it does occur, however, it should be

within the guidelines established at the top. This kind of disciplined implementation will enable the systems developed at any level to be integrated into a hierarchy of systems.

The 1981 Summer Study concluded:

> ... systematic reorganization may well become normal organizational procedure. At the very least, this will require that "change agents" be widely distributed throughout the organization so that people will know where to turn in order to initiate change. It will also require extremely sophisticated computerized planning tools to allow NASA to develop the level of awareness and intelligence needed to guide local and global reorganization.

d. Initial Applications

One possible application of AI is in connection with the NASA Space Systems Technology Model, which is updated annually. The model provides a knowledge base for guiding technology development for future systems and programs. As such, it is a planning document. Beginning in 1983, the knowledge base for the model will be computerized, making possible the application of AI tools in the use of this knowledge base. The conversion of the latter to computer form should be done with this in mind. A natural-language interface between the knowledge base and its human users could be added at an early date. Other AI techniques, such as knowledge-based planning, would probably be introduced later, following experience with such planning systems as DEVISER. It's also possible that the Space Systems Technology Model would be useful in mission planning at, say, one of the NASA centers, where an AI planner would be installed.

Other computerized knowledge bases may be used in the future in conjunction with the NASA Space Systems Technology Model - for example, the NASA Long-Range-Planning Document.

An important benefit that would come from the approach suggested

here is internal consistency of the knowledge base. Merely collecting or computerizing a data base does not make it consistent or validate the items that comprise it. However, if the data base is constructed so that the knowledge is represented in a form subject to "reasoning" of the type performed in an AI expert or planning system, then inconsistencies are more likely to be revealed and corrected.

Once a consistent knowledge-based system has been developed, prospective AI techniques could be used to optimize plans, as well as to replan when either resources (including funding) change or new technology becomes available.

e. Summary

AI will be utilized for all the traditional management functions and at all levels. Natural-language interfaces and expert systems will be used initially Planning, now done largely by hand, will gradually become more automated; the first stage of this process will involve computerizing knowledge bases for existing jobs, such as the NASA Space Systems Technology Model. Knowledge bases will be linked in networks and will in effect become integrated - thus available to a larger number of users. AI will intensify the process of automation by providing natural-language access and the ability to "reason" about the knowledge. The jobs of human managers and planners will change as a consequence and, inevitably, so will the organization.

E. Previously Impractical Missions Enabled by AI

While the applicability of artificial intelligence to present tasks in NASA is important, the new ones that simply could not have been done without AI may turn out to be the most significant.

1. Scaling Up Present Operations and New Missions

The activities that will be possible with AI, that cannot be done without it, are especially interesting. They seem to be of two kinds:

- Operations on a scale that would not be feasible without automation. An example is the telephone system, which would have required more human operators than there are people if automatic switching equipment had not been invented. (Note that a lead time for R&D must always be taken into account.)

- New functions, such as air travel, that would never have become reality without new technology, such as aeronautics.

Our previous discussion included some items of the first category. Among these was the application of AI in the acquisition, storage, transmission, and dissemination of information. In regard to these functions, the 1981 Summer Study concluded that assimilation of the vast quantities of information to be collected by NASA would be impossible without new techniques for massive data storage, data base management, and data processing. A JPL study (McReynolds, 1978) estimated that NASA could save $1.5 billion per year by the year 2000 through a concerted program of AI implementation.

The second category, new functions, involves perhaps somewhat more speculative projections. However, one thing seems obvious: any significant exploration of space will require largely autonomous spacecraft equipped with and supported by artificial intelligence. Space is, after all, an environment of infinite proportions and distances. Remote control of spacecraft by earth-based human controllers will not be possible because of the time required for signals to reach the spacecraft, because of the latter's functional complexity, and because of the large number of spacecraft.

2. MIT ARAMIS Study

The study done by MIT for Marshall Space Flight Center (Miller et al., 1982) explored the potential application of automation, robotics, and machine intelligence (ARAMIS) to space activities, together with their ground-support infrastructure, during 1985-2000. A systematic methodology adopted for the study was designed to cover a wide range of space missions and identify the most promising applications of ARAMIS.

The following representative missions were selected:

- Geostationary Platform (GSP)
- Advanced X-Ray Astrophysics Facility (AXAF)
- Teleoperator Maneuvering System (TMS)
- Space Platform (SP).

These missions were chosen because they span the 1985-2000 period and encompass such a variety of activities, including communications, astronomy, satellite servicing and support, and science and application development. Each space project was broken down into tasks, at five successively more detailed levels - with the smallest tasks, such as "adjust current and voltage," at the most detailed level. Sixty-nine tasks at this most detailed level were chosen for the study. These were then organized into nine types listed below, covering the entire spectrum of tasks that NASA's projects are expected to require during the next 20 years, as follows:

- Power handling
- Checkout
- Mechanical actuation
- Data handing and communication
- Monitoring and control
- Computation

- Decision and planning
- Fault diagnosis and handling
- Sensing

The MIT team then considered what ARAMIS capabilities could be utilized to accomplish the 69 tasks. A classification scheme was developed in which ARAMIS was divided into 6 general areas and subdivided into 28 specific topics, as shown in Table 1. These 28 topics were considered in relation to each of the 69 tasks and, as this process of appraisal unfolded, were further refined into 78 ARAMIS capabilities. Several subjective but systematic methods were used in evaluating the 78 ARAMIS capabilities under consideration.

Table 1

LIST OF ARAMIS "AREAS" AND "TOPICS"

MACHINERY

1. Automatic Machines
2. Programmable Machines
3. Intelligent Machines
4. Manipulators
5. Self-Replication

SENSORS

6. Range and Relative Motion Sensors
7. Directional and Pointing Sensors
8. Tactile Sensors
9. Force and Torque Sensors
10. Imaging Sensors
11. Machine Vision Techniques
12. Other Sensors (Thermal, Chemical, Radiation, etc.)

HUMAN-MACHINE

13. Human-Machine Interfaces
14. Human Augmentation and Tools
15. Teleoperation Techniques
16. Computer-Aided Design

DATA HANDLING

17. Data Transmission Technology
18. Data Storage and Retrieval
19. Data and Command Coding
20. Data Manipulation

COMPUTER INTELLIGENCE

21. Scheduling and Planning
22. Automatic Programming
23. Expert Consulting Systems
24. Deductive Techniques (Theorem Proving)
25. Computer Architecture

FAULT DETECTION AND HANDLING

26. Reliability and Fault Tolerance
27. Status Monitoring and Failure Diagnoses
28. Reconfiguration and Fault Recovery

One approach focused upon the relationships among the ARAMIS capabilities, in view of the fact that the development of some ARAMIS capabilities depends upon the development of others. These relationships were described graphically in a "technology tree."

Another method established seven decision criteria:

- Time to complete the task
- Maintenance
- Nonrecurring cost
- Recurring cost
- Failure proneness
- Useful life
- Developmental Risk.

Then, for each of the 69 tasks, these 7 decision criteria were applied. Points were assigned subjectively for each criterion, on a scale of 1 to 5. The points were allotted to each of the candidate ARAMIS capabilities for that task, one prospect being the existing technology that would be used to accomplish the task. Sixty-nine Decision Criteria Comparison Charts were developed, one for each task.

Using the technology trees and the 69 charts of decision criteria values, the group then systematically evaluated the candidate ARAMIS capabilities for each type of task. Some of the conclusions that relate to AI applications are as follows:

- Power handling - An onboard adaptive control system capable of modifying its own programming.
- Checkout - A computer simulation of mission sequences, either prior to launch as part of spacecraft verification, or after launch to support mission decisions or failure diagnosis.
- Mechanical actuation - A dexterous manipulator under human control.

- Data handling and communication - Fault-tolerant software.

- Monitoring and control - An onboard adaptive control system.

- Computation - For logical operations and evaluations, an expert system with human supervision and a learning expert system with internal simulation.

- Decision and planning - Computer modeling and simulation.

- Fault diagnosis and handling - An expert system with human supervision. "The study group feels that expert systems may become not only desirable but necessary in future spacecraft missions. The traditional philosophy is to anticipate all possible one-point and two-point failure modes during the design process, and to design either safeguards or recovery systems to deal with possible problems. However, as spacecraft complexity increases, the prediction of all such failure modes and effects becomes combinatorially enormous."

The MIT group concludes that AI may be necessary for even planned space missions and that the most critical application may be in the area of fault diagnosis and handling.

3. General-Purpose Robot Explorer

The 1980 Summer Study studied three future space missions that would require AI, including a general-purpose robot explorer spacecraft. They considered in detail a demonstration mission to Titan, largest of Saturn's moons. Titan was chosen in part because it lies far enough from earth to preclude direct intensive study of the planet from terrestrial observation facilities or easy teleoperator control, yet is near enough for system monitoring and human intervention. The target launch date for the Titan demonstration was set for 2000, with five years on site. The team that studied this mission concluded that the most important single technological factor in making automated space-exploration missions of the future feasible is the potential capacity of artificial intelligence for learning in and about new environments and for generating scientific hypotheses.

4. Automated Manufacturing System in Space

The 1980 Summer Study looked at a permanent, highly automated, general-purpose, space manufacturing facility that would evolve into a self-contained production cell, independent of material supply from earth. Several basic "starting kits" were described in the study, which also devoted some of its attention to the concept of extracting raw materials from the moon and asteroids. The study concluded that a manufacturing facility in space would require that system functions be taken over by AI and autonomous robots, thus accelerating a trend already evident in industrial robotics.

A more immediate use of AI for automated assembly in space (e.g., in geosynchronous or low earth orbit) with materials supplied from earth is a very attractive possibility. Semiautonomous robots in orbit would enable many tasks, such as assembly of large antennas, space station construction, repair and maintenance, construction of interplanetary space vehicles, and repair of satellites. These tasks can be accomplished without the high cost and risk of manned assembly, especially as the more difficult portions of the robot's tasks can be directed from earth.

There are several hardware considerations that make it simpler to design and build robots to accomplish these missions than to design and build equivalent terrestrial robots. For example, in space the power-to-weight ratio of a robot is not important because this only affects the speed at which objects can be moved (which should be kept low for safety reasons anyway). Although production of precision robots with high power-to-weight ratios is the main design problem for earthbound robots, very "weak" robots can perform quite well in space. Furthermore, space forms a perfect background for existing vision systems because of its pitch blackness — except for the stars and planets, which

provide excellent reference points for directing and orienting the camera. The vacuum of space is also uncluttered and clean and is a perfect insulator, allowing metal-to-metal contacts to be detected by changed conductivity without the possibility of stray conduction paths. The significance of these special conditions is that NASA could gain substantial capability to do manufacturing, repair, and assembly in space - at low cost and in the near future - by using specially designed robots. These robots and their sensors would be controlled by AI planning systems that could either be directed from earth to any level of detail or could modify their own behavior, as necessary, to achieve the assigned objectives.

5. Self-Replicating Systems

The 1980 Summer Study also proposed an automated, multiproduct, remotely controlled, reprogrammable lunar manufacturing facility capable of constructing duplicates of itself that would themselves be capable of further replication. The required technology is based, in part, on CAD, CAM, CAT, and robotics. The study team noted that, when the techniques CAD, CAM, CAT, and robotics are used to produce components of their systems, a high level of automation that Albus (1976) has called "superautomation" is achieved. The self-replicating lunar factory is illustrated in Figure 5.

Self-replicating systems have been studied at Marshall Space Flight Center, starting with a theoretical foundation from which the engineering concepts for self-replicating systems have been developed. One concept is for a "universal constructor," a device that can build any machine if provided with proper instructions (von Tiesenhausen and Darbro, 1980).

A number of opportunities for NASA in space were considered at Woods Hole. While not dealing specifically with artificial intelligence, The Innovation Study assumed advances in AI and robotics as a basis for most of the advanced concepts included in its report.

ORIGINAL PAGE IS
OF POOR QUALITY

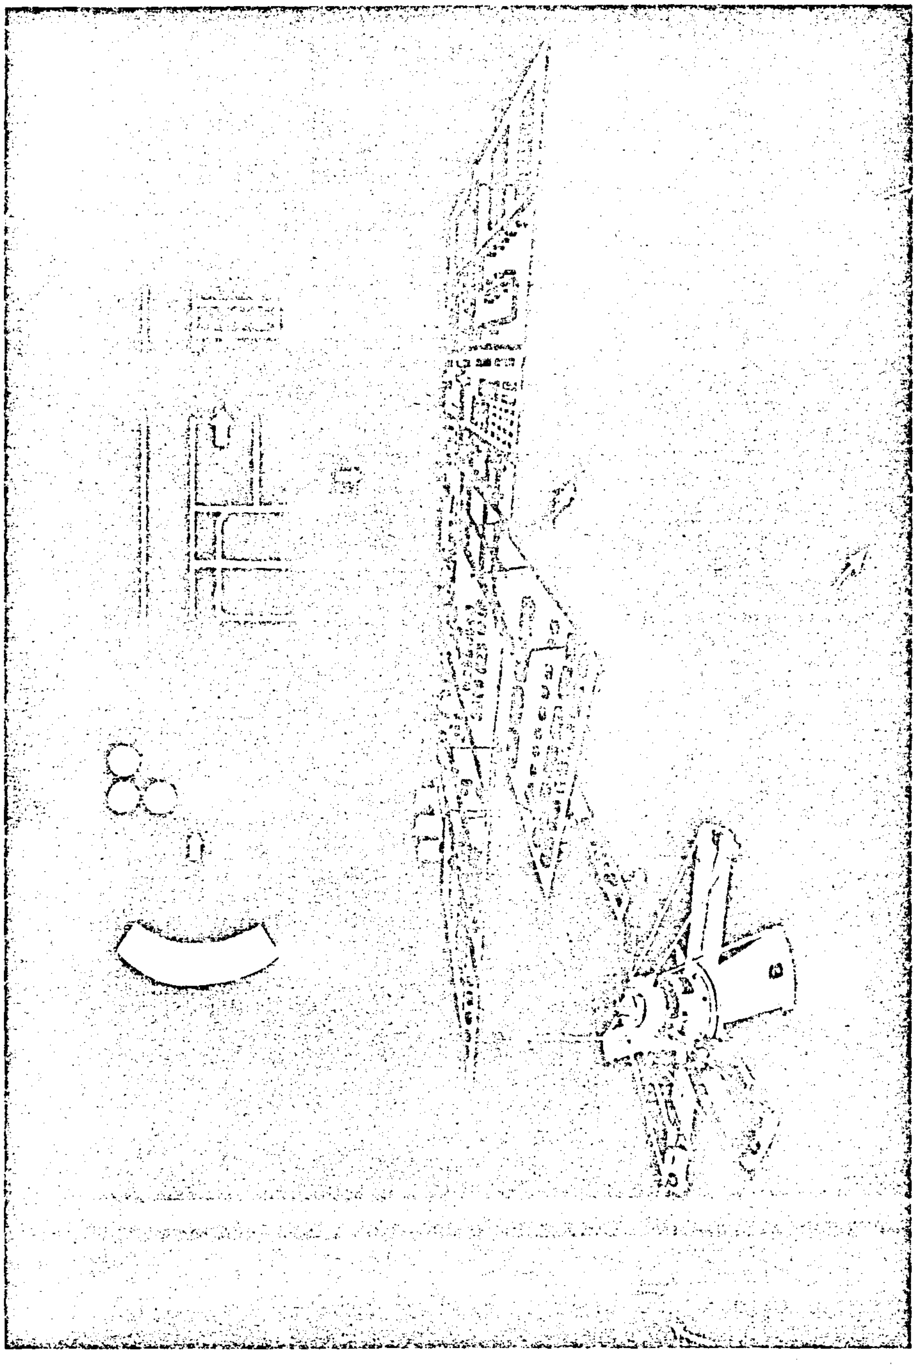

FIGURE 5 SELF-REPLICATING LUNAR FACTORY

As regards the subject of self-replicating systems, the potential difference in output between linear and exponential systems was recognized as phenomenal.

6. Other Applications in Space

Because ground-space communications are vulnerable and will be increasingly overloaded, onboard AI systems to operate, diagnose, plan, and replan for such future missions as the space station are essential, especially for time-critical responses. In addition, AI systems could help configure onboard equipment and instruments (e.g., the R1 system, used at Digital Equipment Corporation for automatically generating VAX configurations in detail), schedule crew tasks, and provide intelligent interfaces between subsystems. These interfaces could be between the crew and station equipment, between the ground and station systems, or between station systems themselves.

Space-station robot operations, especially those involving high-speed local feedback, will have to be run by local AI expert systems for planning and replanning. This includes such operations as space assembly and external station maintenance.

Because space is a unique environment, NASA will have to develop its own systems for these applications. Some of the necessary AI technology can be expected to come from normal development in the field, but the special conditions of space will require specialized knowledge bases and robot control.

The foregoing considerations highlight the fact that an AI system possessing some ability to plan, monitor, diagnose, and replan, or at least to advise and assist a person to do these tasks, is essential for time-critical operations and extremely useful for normal operations.

As far as deep-space missions are concerned, fully autonomous AI systems that can plan, monitor, diagnose, and replan are essential because of earth-based communication delays. Such an AI system must be able to model itself and all its subsystems, since its ability to repair and control itself without effective contact with earth requires that it have such a model. This self-modeling ability is especially important for autonomous missions in deep space, but it also has obvious advantages for performing similar tasks on earth (e.g., mining deep underground). Although NASA will benefit from general robotics/AI research, its unique deep-space requirements (especially the need for very high reliability and the extreme degree of autonomy) will inevitably impel NASA to do AI research in the direction of its special needs.

A survey of the future of robotics and automation in space was published in 1978 by Heer. The findings, based on the 1978 NASA Space Systems Technology Model, were summarized in the the chart shown in Table 2. The survey concluded by stating that "... a new level of robotics and automation should catalyze some missions, but the technology will mainly be developed and applied to reduce operational costs for all missions."

Concepts that are even more advanced, such as space settlement, presuppose a highly developed AI technology (Johnson and Holbrow, 1977). The degree of automation and flawless maintenance that will be required will necessitate intelligent systems that will, most assuredly, incorporate artificial intelligence.

The search for extraterrestrial intelligence (SETI) can be aided by AI in at least two ways. Since SETI is a difficult search problem and heuristic search is a prime subject of AI research, a direct application is possible for automation of the search process (Morrison et al., 1977). Moreover, an understanding of intelligence and models of intelligence are essential for SETI - and it is precisely toward the achievement of those goals that AI research is directed. ("The CETI Program," 1974).

ORIGINAL PAGE IS
OF POOR QUALITY

Table 2

ESTIMATES OF THE TECHNOLOGY DEVELOPMENT EFFORTS TO AUTOMATE SYSTEM FUNCTIONS

SYSTEM FUNCTIONS

GROUND OPERATIONS: DATA INTERPRETATION; DATA DISTRIBUTION; ARCHIVING; SEQUENCING; SITE SELECTION; SPACECRAFT MONITORING; SIMULATION

ON-BOARD SPACECRAFT OPERATIONS: POINTING AND CONTROLLING; GUIDANCE AND NAVIGATION; RENDEZVOUS AND STATION KEEP.; DOCKING; LANDING AND ASCENT; SURFACE TRAVERSING; SCIENCE DATA COLLECTING; SERVICE DATA COLLECTING; DATA PROCESSING; PLANNING DECIDING; RESOURCE CONTROL; FAULT DETECTION, REPAIR; SHAPE CONTROL

IN-SPACE HANDLING(2): DEPLOYMENT; TRANSFER; ASSEMBLY AND JOINING; INSPECTION; SERVICING; MAINTENANCE; REPAIR; MONITORING; RETRIEVING; RESCUE; PROCESSING; MANUFACTURING; LUNAR MINING

MISSION CATEGORIES		
EXPLORATION OF SPACE	GALILEO-JUPITER ORBITER PROBE	1982
	SOLAR POLAR MISSION	1983
	VENUS ORBITAL IMAGING RADAR	1983
	SPACELAB INSTRUMENT PROGRAM	1984
	X-RAY OBSERVATORY	1985
	SATURN ORBITER DUAL PROBE	1987
	MARS SAMPLE RETURN	1988
	MOBILE LUNAR SURFACE SURVEY(1)	1989
	EARTH ORB. SOLAR OBSERVATORY	1992
	ASTROPH. SPACE LABORATORY	1993
	ATMOSPHERIC PHYSICS LAB.	1994
	SPACE-BASED RADIO TELESCOPE	1995
	AUTOMATED PLANETARY STATION	2000
GLOBAL SERVICES	SEASAT FOLLOW-ON	1982
	TIROS-O	1984
	SOIL MOISTURE SATELLITE	1985
	STORMSAT	1985
	GLOBAL COMMUNICATIONS	1987
	GLOBAL CROP FORECASTING	1988
	HIGH RESOLUTION SEA SURVEY	1988
	DISASTER WARNING SATELLITE	1989
	EARTH ENERGY BUDGET MONITOR	1992
	LGE-SCALE ALL-WEATHER SURVEY	1992
	GEOLOGICAL MAPPING	1994
	GLOBAL NAVIGATION	1996
UTILIZATION OF SPACE	SPACE MANUFACTURING MODULE	1985
	SPACE HEALTH CARE SYSTEM	1986
	LUNAR PRECURSOR PROCESSOR	1990
	NUCLEAR WASTE DISPOSAL	1995
	TELEOPERATOR VEHICLE SYSTEM	1996
	ROBOT VEHICLE EARTH ORBIT	1997
	LUNAR BASE	1998
	SPACE STATION	2000
	SATELLITE POWER SYSTEM	2000
TRANSPORTATION SYSTEMS	ORBITAL TRANSFER VEHICLE	1990
	HIGH ENERGY OTV (PLANET)	1992
	PRIORITY LAUNCH VEHICLE	1994
	HEAVY-LIFT LAUNCH VEHICLE	1998

KEY THE AUTOMATION OF THE IDENTIFIED SYSTEM FUNCTIONS REQUIRES

/ INTEGRATION OF EXISTING TECHNOLOGY

X MODERATE ADDITIONAL DEVELOPMENTS

EXTENSIVE TECHNOLOGY DEVELOPMENTS

MAJOR TECHNOLOGY DEVELOPMENTS

MAJOR TECHNOLOGY DEVELOPMENTS WITH UNCERTAIN OUTCOME

NOTE: EACH ENTRY REPRESENTS THE RELATIVE COLLECTIVE LEVEL OF EFFORT TO ACCOMPLISH THE FUNCTION FOR THE MISSIONS AS DESCRIBED IN THE NASA OAST SPACE SYSTEMS TECHNOLOGY MODEL, MARCH 1978.

1) THE LUNAR ROVERS OF THIS PROGRAM WILL BE DEVELOPED WITH IN-SPACE HANDLING CAPABILITIES AND WILL SUPPORT THE LUNAR PRECURSOR PROCESSOR (1990) AND THE LUNAR BASE (1998).

2) HANDLING FUNCTIONS ARE GENERALLY ASSOCIATED WITH MOBILITY UNITS, MANIPULATIVE DEVICES OR TOOLS REQUIRING CONTROL OF ACTUATORS

7. <u>Applications in Aeronautics</u>

In addition to space missions, other future systems of concern to NASA depend on advances in AI and robotics. For example, the MIT ARAMIS study concluded that AI will be necessary to obtain the degree of reliability needed for future space missions. For the same reasons, AI will be an absolute prerequisite for future, high-performance, fly-by-wire aircraft that humans alone cannot control. Such aircraft are desirable because they can be more energy-efficient, but the practically instantaneous response time required by the control system is too short for human pilots.

Another possible application in aeronautics is in computational fluid dynamics. At the present time, experienced aeronautical engineers (human experts) decide where to use either fine or coarse gridding, a decision that can have a substantial effect on the amount of computation required. Their decision is based on their best judgment as to where laminar flow can be expected, where turbulence may occur, etc. The Applied Computational Dynamics Branch at Ames Research Center is now looking at possible AI techniques that can be used to perform this function.

8. <u>Automatic Programming</u>

One of the most rewarding potential gains from AI, however, will likely come from the field of AI research known as automatic programming (AP). AP systems would write their own applications programs according to specifications that describe exactly what the program is intended to achieve, thereby reducing the time, effort, and expertise required in the production process. Knowledge-based AP systems differ significantly from the so-called automatic program generators currently being marketed. These rudimentary systems simply select program features from a menu in much the way an individual might select factory options on an automobile. A true AP system, on the other hand, would be analogous to

feeding into a computer the basic specifications for an automobile, such as maximum speed, number of passengers, etc., and having an automobile automatically designed in conformance with those specifications. A major example of experimental work in this field is the CHI system developed by Cordell Green at the Kestrel Institute (Barr and Feigenbaum, 1982, pp. 326-335). CHI is an interactive, knowledge-based programming environment that emphasizes the use of a very-high-level programming language, V, both for programs and for program knowledge.

Offshoots of AP are program transformation and program verification. Program transformation is the process of systematically rewriting a program without altering its external behavior, generally for the purpose of improving its efficiency. Program verification is the process of proving that a program satisfies a given specification. These two offshoots of AP are, perhaps, even more important than the more generalized approach of program synthesis, since maintenance tasks (including editing, debugging, and modification) are estimated to be increasing at an alarming 22 percent per year. Maintenance currently consumes up to 35 percent of the average programmer's time and up to 70 percent of most programming budgets.

This Page Intentionally Left Blank

VI CATALOG OF AI RESEARCH AND DEVELOPMENT

The following is a summary of the main areas of research and development in artificial intelligence. AI research activity is split between basic research on fundamental issues, such as knowledge representation, and research on specific areas of application, such as vision, planning, etc. Of course, all application areas utilize advances in basic AI; consequently, to reap the full benefit therefrom, NASA should support basic AI research to encourage its growth and to help develop an expanding reservoir of professional expertise in this important area. The fact that there are problems connected with AI applications that are unique to NASA's mission makes the support of AI research and development necessary and, as solutions are found, ultimately rewarding.

Most AI research in the United States is conducted in universities, with the remainder predominantly in research institutes. These centers each support research in nearly every area of AI described below. The leading AI establishments (in alphabetical order) are

Carnegie-Mellon University (CMU)
Massachusetts Institute of Technology (MIT)
The Mitre Corporation
The Rand Corporation
Rutgers University
SRI International (SRI)
Stanford University
University of Rochester
Yale University

In addition, many computer science departments of other universities support AI research on a lesser scale, while some large corporations (e.g., Hewlett-Packard and Schlumberger/Fairchild) and AI start-up companies (e.g., Teknowledge) support AI research and development, usually with a strong applications orientation. All the institutions mentioned above are involved in the entire spectrum of AI research, although the emphasis varies from one establishment to another; consequently, they are not listed separately under each of the research headings below. The following sections outline the basic areas of AI research and development, with particular stress on their relationship to NASA.

A. Basic Research

1. Heuristic-Search Theory

In the early days of AI research, it was thought that sufficiently fast and efficient search procedures would enable the solution of difficult problems. This expectation remained unfulfilled, however, as it became apparent that without sufficient knowledge to guide the search, the system would waste its time investigating unproductive branches. Since this revelation, AI research has investigated heuristic search guidance techniques, with considerable success. There is now a large body of both theory and actual results (for example, see the recent article by Pearl, 1981).

2. Knowledge Representation and Acquisition

During the 1970s, it was realized that AI research should be directed largely at general knowledge representation and the methods of acquiring and using this knowledge. In particular, representations for

time, uncertain and incomplete knowledge, processes and actions, and knowledge about other agent's knowledge, have been investigated.

There is divided opinion in the AI community regarding the best framework for building particular representations. On one side are those who have developed high-level representations or representation languages with an intended intuitive interpretation, while, on the other, are those who build more formal representations in mathematical logic. The first group had initial success with such representations as semantic nets and causal inference networks, but these methods later encountered difficulties of interpretation because the meaning of these representations had never been precisely defined. Although logic had always had a clear meaning, it suffered from computational difficulties. The development of computer languages based on logic, such as PROLOG, have removed many of these computational barriers.

Regardless of which type of representational framework is used, any nontrivial domain requires that the appropriate concepts and relationships be discovered and that an extensive knowledge base using this framework be built. Even when the appropriate language is obvious a large investment is necessary to build a working debugged system.

3. Commonsense Reasoning and Problem Solving

This area of AI research is concerned with developing representations and inference procedures for real-world reasoning, including the use of uncertain and incomplete information. Such representations should reflect the level of detail necessary for the particular application. For example, it is often unnecessary to solve the heat flow equations in detail to deduce that an absorbent surface exposed to the sun may lead to overheating. Such rough reasoning, sometimes referred to as "naive physics" or "commonsense reasoning," is often sufficient to solve problems without detailed calculations.

The development of representations that capture the essential properties of a particular domain without becoming excessively complex can be a surprisingly difficult task. For example, one area of naive physics that has been extensively investigated is the representation of liquids. This domain raises such questions as: Is a lake still the same lake after all the original water in it has been replaced? and How does one represent flow and overflow in different situations? Because of the uniqueness of the space environment, NASA will have to invest considerable time and money developing suitable representations for every major application area.

4. AI Languages

Because of the complexity of AI systems, there has been a trend to develop better programming tools so that researchers can concentrate on the application rather than the software support. This effort has resulted in advanced (high-level) languages that make symbolic representation and reasoning easy, and in programming-support environments that automatically take care of such details as file handling, online debugging and editing, etc. One of the earliest symbolic languages was LISP, introduced in 1958 (McCarthy et al., 1965). It has proved so useful that it is still the main language of AI research more than twenty years after its introduction. Although the basic structure of the language has remained unchanged during this period, a powerful support environment has been developed (e.g., INTERLISP), that relieves the programmer of much of the software development burden.

A more recent AI language that is gaining popularity is PROLOG. This is a rule-based language that offers certain advantages over LISP in some applications. One advantage is that PROLOG allows a procedure to be used in multiple ways, depending on which arguments are bound when it is called. In addition, PROLOG includes a data base system as part of the language, thus avoiding the need to write one in LISP. PROLOG can be regarded as the next step toward even higher-level languages that

allow the user/programmer to specify what is to be done then leave the details of execution to the language system. In PROLOG there is a clear separation between the logic of the domain (expressed as facts and rules) and the control that dictates how this specification will be used to answer particular queries. In the future, even less control information will be required from the user.

As AI research continues, enhancements of the basic AI languages can be expected mainly by continuing to build the programming support environment An increase in speed can also be expected, aided by the introduction of dedicated, special-purpose hardware. Examples are the LISP and PROLOG machines of the Japanese fifth generation computer project (Feigenbaum and McCorduck, 1983). These software/hardware tools will facilitate the ruture development of expert systems, as will the even higher-level languages that can be expected to emerge.

B AI Applications

1. Expert Systems

Expert systems are programs that capture an expert's expertise in a particular area, for subsequent use by a nonexpert. Such programs differ from normal special-purpose programs, such as those for spacecraft navigation, in that the embodied expertise is usually approximate and uncertain. If there is a well-developed theory for the domain of interest, this should be embedded in a special-purpose program, by using appropriate software tools. However, when the only knowledge about a domain is in the form of "rules of thumb," uncertain associations, and only partially understood relationships, an expert system can usually be built to utilize such information. For a survey of AI research in expert systems see Gevarter (May 1982).

Some of the domains to which expert systems have been applied, together with the corresponding programs, are listed below:

Medicine - The MYCIN system (infectious diagnosis and therapy recommendation - Shortliffe, 1976), and
EXPERT (general diagnosis - Weiss et al., 1979).

Geology - The PROSPECTOR system (Duda, 1980), which has been used to locate a molybdenum deposit, and is currently being employed in a survey of selected areas for possible useful deposits.

Signal Analysis - The HASP/SIAP program (Nii et al., 1982), which analyzes underwater acoustic signals to detect ships and other objects despite a high level of background noise.

Fault Diagnosis - The EL program (Sussman, 1977), can simulate the operation of an electrical circuit and thus deduces the possible cause of a failure; it can even suggest specific tests to track down a fault in the system.

The above applications are all characterized by a varying degree of uncertainty in the information supplied, as well as in the rules used. These expert systems are still able to perform well under conditions of uncertainty because they use methods of inference similar to those employed by an expert under the same circumstances. Conventional computer science is unable to produce programs that handle uncertainty with anything like the ability of expert systems, because they process all information with the same algorithm rather than by adapting to the information available.

The general categories of tasks that expert systems have been applied to can be broken down as follows:

a Interpretation/Diagnosis - This category of expert systems includes all those that can accept data from the user about a particular case and when sufficient information has been received, return a diagnosis or interpretation of that case. Examples include mass spectroscope data interpretation (DENDRAL - Buchanan and Feigenbaum 1978) and most of the medical expert systems (MYCIN, EXPERT, etc.).

b Design Systems - These are expert systems that may be given particular information and constraints and are required to produce an output that satisfies the given design criteria. They include, for example, planning systems that are required to produce a plan in a

particular situation, chemical synthesis programs that generate a procedure for synthesizing a particular molecule, and music composition programs whose output conforms to a prescribed style and form.

c. Prediction and Induction Systems - These systems accept data and look for patterns or other forms of order. When such patterns are found, they can be combined with information about a particular case to predict the most likely outcome. In most expert systems, the discovery of patterns is performed by the expert prior to knowledge transfer; in many domains, however, either there are no experts or the patterns are very weak and deeply buried in the data. In such cases, an inductive expert system is essential. An example of an inductive system is INDUCE (Michalski, 1980), which inferred the relationship between symptoms and disease in soybeans.

d. Monitoring and Control Systems - These systems receive specific online data from the sensors on the object being monitored/controlled. This data is rapidly interpreted by the expert system and the appropriate responses generated. In a monitoring expert system, particular ill-defined situations are represented that will trigger specified alarms if they are ever detected. Control systems, on the other hand, can initiate complex commands to try to bring the system that is being controlled back within operational parameters. Often the generation of an appropriate response will require simulation of the expected effects of possible actions on the controlled system. Such expert monitoring/control systems differ from conventional computer controlled feedback systems in that they can respond to complex situations a programmer may not have thought of by applying domain expertise, in the form of causal models, to diagnose and correct whatever problem or problems may have arisen.

2. Features of Expert Systems

In each of the types of expert systems described above, the following features are usually found:

a. An Explanation System - The system should always be able to retrace its reasoning in a particular case and explain what it did at each step and why. This explanatory capability enables the user to accept or reject the system's conclusions if he disagrees with its reasoning, and aids the expert in debugging the system. Some expert systems provide a "sensitivity analysis"; this information tells the user what input data are most important to a system in reaching its conclusions.

b. Modularity - One of the main reasons for preferring AI expert systems over conventionally programmed systems with a great deal of built-in expertise is that, in AI systems, it is easy to localize errors to particular facts and then to modify the latter without disrupting the program. This capability stems directly from the "rule-based" architecture of AI systems that store their knowledge in the form of modular facts and rules so that a particular modification will not alter others. These facts and rules are interpreted by a domain-independent program so that such modifications will not affect the interpretive program. Such data-driven programs are also more transparent to the expert and user.

c. Use of Models - There is a division between "shallow" and "deep" expert systems. In a shallow system, the surface symptoms (features) are related directly to the corresponding diagnosis according to the rules provided by the expert. A deep system, on the other hand, relates a case to its own model of the situation, using the model to propagate possible effects and deduce possible corrective actions. In a shallow system the expert has incorporated his understanding of the domain in the form of rules relating the symptoms to the diagnosis, thereby rendering the system unable to reason from first principles about the domain. Deep systems are obviously harder to build because of the difficulty of constructing a computer model, but they are much more flexible and can be applied to a much wider range of potential situations.

The main difficulties that must be surmounted in developing expert systems are, firstly, finding general methods for combining diverse pieces of uncertain information to reach a probabilistic conclusion and, secondly obtaining, representing, and debugging expert knowledge about a particular domain. Typically this last step may take several man-years of effort. This investment is usually worthwhile, since the resulting expert system can then be used repeatedly - whenever and wherever it is needed Furthermore, methods are currently being developed for dealing with these problems that should reduce the time it takes to build new systems. Basic research in expert systems is continuing at most AI research centers and companies already exist (e.g., Teknowledge) to aid customers in developing their own expert systems.

3. Planning

Considerable AI research has been done on planning systems (an important type of expert system). This has resulted in many advances, such as techniques for representing time (e.g., DEVISER - Vere, 1982) and geometric modeling in such programs as ACRONYM (Brooks, 1981). These systems have been applied to such tasks as controlling autonomous robots, assembly planning military, and space missions. Any nontrivial planning task utilizes a large amount of domain-specific knowledge to guide the planner at every step. The major use of this knowledge is to simulate the relevant portion of the world in which the plan is to operate so that the planner, prior to execution, can detect any potential difficulties in its current plan. In addition, considerable heuristic knowledge is required to guide the planner's choices while building the plan, so as to ensure efficient exploration of the entire spectrum of possibilities.

If a planner is designed to be interactive - that is, if the user can make choices, and ask the system for information (and vice versa) - the system then resembles the current generation of expert systems. An example of this type of system is KNOBS, developed by the Mitre

Corporation to be used for Air Force mission planning. However, a planner can be designed to be completely autonomous (e.g., DEVISER), an essential requirement for controlling an autonomous vehicle or for semiautonomous operations in space.

4. Theorem Proving

Theorem-proving research has been going on since the inception of AI because a successful theorem prover would allow the automatic deduction of consequences in any domain that could be formally represented. Despite initial successes, theorem proving has failed to match its promise, leading to a decline in theorem-proving research. The main reasons for failure have been the difficulty of representing interesting domains and the computational difficulties in controlling the search space without domain specific heuristic guidance. Subsequent research has emphasized knowledge representation and methods of controlling deductions. A useful by-product of this research is the language PROLOG, which allows the user to control the deduction process.

5. Vision

One of the worst bottlenecks in using computers is the problem of transferring information to the machine. A keyboard is the usual means of information transfer, but, given the dictum that a picture is worth a thousand words, it is clear that the computer could understand the world much better if it were possible to use a camera to supply visual information that the computer could understand. The advantages of vision for computers that interact with the real world has been long realized in AI, and there is a considerable body of research in this area. The major problems have to do with interpretation of the resulting digitized image rather than with the hardware needed to obtain such images. For a survey of AI research in computer vision see Gevarter (September 1982).

Research on computer vision is still in its infancy, but already there are commercially available AI-based vision systems that are routinely being used in automated factories for such tasks as inspection and recognition, as well as guiding robots in assembly tasks. These achievements are possible only because the work environment can be highly constrained, so that, for example, the camera is in a known position relative to the work-bench and the lighting is optimally arranged.

The main approach in current AI vision research is to try to relate the image to an unknown world model by means of such clues as edges, regions, color, and shading contours to suggest plausible three-dimensional world models. These models are then further refined by knowledge about real-world constraints and reference back to the image. The ability of such experimental programs to recognize real images has improved over the last decade, but is still far from achieving general scene recognition. Even when an experimental vision system does recognize objects in a scene, the process is generally very slow - on the order of minutes. However, this time limitation should not be considered important, since the speed of computers can be expected to improve substantially over the next decade and most vision procedures are designed so that they can execute in parallel.

Of major significance to NASA is the fact that space offers an ideal environment for the primitive vision systems that are currently available. This is because space itself is a near perfect black background. The few features in space (stars, planets, sun, and moon) make excellent markers, since their precise positions are known in advance and so they can be used to accurately locate the position of some object (e.g., a structural member) in space. NASA should therefore start its research on robotic assembly in space beginning with currently available vision systems. As regards support of such research, SRI International's Artificial Intelligence Center maintains a vision testbed that contains most of the advanced experimental AI vision systems developed in the U.S.

5. Natural Language

Because of the high cost of training programmers and time such training requires, it would be highly desirable if the users could communicate with the system in English (or some other natural language) rather than in "computerese." Unfortunately, this is a very difficult task when the target is a system with the English capability of an adult speaker, in view of the inherent ambiguity and context-sensitive character of natural languages. For example, a major stumbling block in natural-language research is the problem of resolving reference (known as anaphora), particularly pronominal reference. To illustrate how difficult this is, consider the following pair of sentences:

> The councilmen refused the women a permit to demonstrate because they feared violence.

and

> The councilmen refused the women a permit to demonstrate because they advocated violence.

These sentences differ by only one word, but a human has no difficulty in seeing the shift in reference of the pronoun "they." The reasons for this shift involve a complex combination of knowledge of politics and society that a computer would also have to have before it could be expected to disambiguate such sentences. The representation and use of general commonsense knowledge for understanding is a major topic of AI research.

Another major difficulty in computer understanding of natural language is finding the intended meaning of abbreviated sentences (technically referred to as ellipsis). For example, having just been told to put something down, a person might respond "Where?", which is short for "Where should I put it?." To understand the meaning of such a response, a computer would have to understand what is happening, what the goals of the various actors in this scenario are, and so on.

There are also problems in understanding the ambiguous uses of such words (quantifiers) as each, every, some, any, all, most, etc. For example, does "Will everyone on board please throw a piece of furniture overboard" mean that everyone should pick up the same piece of furniture or a different piece per person? Resolving such quantifier ambiguity requires understanding the situation and employing considerable world knowledge.

It should be clear from the above difficulties that there is no simple solution to the problem of fully understanding natural language in man-machine interaction, although AI research has made considerable progress in this direction. Fortunately, experience with natural-language interfaces to, say, a data base has shown that, in a restricted domain, a degree of ambiguity can be resolved because of the limited number of interpretations that can be put on a particular sentence. In such cases, experimental and commercial natural-language interfaces have been constructed that enable naive users to interact usefully with the system. Such systems avoid the tricky disambiguation problems by asking the user to rephrase the input. Users soon learn how to phrase inputs so that the system understands them. Besides the few commercially available natural-language systems, considerable work is being done in this area by all the major AI research establishments. In fact, there is a well-established and active professional organization (the Association for Computational Linguistics) that holds regular conferences to coordinate, publicize, and promote natural-language research.

6. Robotics

There has been a lot of interest in this subject in recent years, as reflected by a growing number of journals and professional publications devoted to robotics and automation in general. Principally because of an increase in industrial applications, a large amount of

robot hardware has been developed that is a considerable improvement in price and performance as compared with earlier models. Robots are either arms on a fixed base (called manipulators) or mobile carts with some sort of pick-up-and-place capability. A modern manipulator can achieve high accuracy and repeatability in performing actions and can lift up to one-tenth of its own weight without sacrificing performance. This strength-to-weight ratio is of great importance in terrestrial applications (e.g., in servicing the shuttle on the ground), but it is of little consequence in space, where this ratio will only affect the speed at which things can be moved. For a survey of AI research in robotics see Gevarter (March 1982).

As in computer applications, the main difficulty in the use of robots is not the cost or capabilities of the hardware, but the cost of programming them to achieve the desired goals. Even in a factory, which is a highly controlled environment, the major bottleneck in setting up a new assembly is in programming all the automated systems (including the robots) to work in proper unison. This is not surprising, since even an artificial environment, such as a factory, has far more elements that can go wrong than most programs typically deal with. The problem of programming an autonomous robot in an unstructured and possibly unknown environment is clearly beyond current conventional programming methodologies.

AI was applied successfully to the problem of controlling an autonomous robot as long ago as 1972 (STRIPS - Fikes et al., 1972); further research since then has added to and greatly improved the techniques for intelligent robot control. Most of this research comes under the heading of planning expert systems, since producing a plan for a robot to execute (that will then achieve the given goals) is just a particular form of planning.

There are numerous difficulties in automatically producing plans for a robot or group of robots that are sufficiently accurate and robust to be useful. These difficulties include the following:

a. Use of Sensors (e.g., cameras, force-detectors, and contact switches) - The main difficulty in planning with sensors is to be able to model the sensor and the information it can provide so that the plan can include calls to use specific sensors at particular stages in the plan as needed. It is still an AI research problem to produce plans that use sensor-provided information to guide execution along alternative paths and detect when error conditions occur, but considerable progress has been made toward this objective. If the United States is to maintain a lead in robotics technology, a major investment in robotics planning research is the best strategy. Moreover, the end result would be very useful to NASA in both terrestrial applications (e.g., shuttle servicing) and in space.

b Cooperating Robots - There are many problems entailed in representing the actions of multiple robots that are cooperating to perform some task. If all the robots are controlled by a single controller, the main difficulty in producing plans for such a system is to ensure that there will be no harmful interactions among the actions that different robots are performing simultaneously. This interaction detection can be performed by an appropriate world simulation.

If there is a separate planning system on board each cooperating robot, every planner has to model the other planners' goals and possess the knowledge to predict what the other robots will do. There are many difficulties in representing such knowledge fully, e.g., such infinite regressing as "He thinks that I think that he thinks that ..." and problems of deadlock when one robot is waiting for another to finish while the second robot is waiting for the first to finish, etc. As there is considerable AI research in progress on representational problems of this kind, reasonable working systems can be expected in the intermediate future. Research on such problems is essential if projects like the lunar self-replicating automatic manufacturing and assembly station are to be given serious attention. At the same time, more immediate problems, such as assembly in LEO with multiple robots, must be dealt with and resolved The main robotics research centers are MIT, Stanford, CMU

(particularly the new robotics center), SRI, and NBS, while most large corporations (e.g., IBM, GE, DEC, Boeing, etc.) have their own robotics research and applications groups.

7. Automatic Programming

To a limited degree current compilers already do "automatic programming." They are given a description of what a program is to do in a high-level language, then write a machine-code program to do it. Automatic programming in AI can be viewed as a "super compiler" - i.e., a program that can accept a very-high-level description of what the program is to accomplish and produce a working program to do it. The high-level description might be in a precise formal language such as logic or set theory, or it might be a "loose" English description that requires further dialogue between the system and the user to resolve any ambiguities.

The task of automatically writing a program to achieve a stated result is closely related to the task of proving that a given program actually achieves a stated result. This proof task is called program verification Many automatic programming systems produce a verification of the output program as an important side benefit.

In one form of automatic programming, the user describes the intended behavior of the target program by giving examples of input and its corresponding output. Since this method of instruction is inherently ambiguous, it is possible that programs thus produced will not necessarily do what the user intended in all circumstances, in contrast to the results of formal specifications. Such automatic programming-by-example systems are exhibiting a type of learning.

There has been considerable progress to date in the area of automatic programming and programs of moderate complexity have been synthesized. Given NASA's already large investment in customized software

and its anticipated future growth, the emergence of practic. automatic-programming systems will help make both new softw... ment and the reimplementation of existing software much more ... In addition, the existence of such systems will enable comp... tions to be developed that are too difficult or costly today, ... assure greater reliability in the end product. Consequently, ... area of AI in which a NASA investment can yield an enormous re...

8. Learning

Learning is a generic term that applies to all AI progra... improve their behavior as a result of experience. Give this ve... definition, it is not surprising that there are many differ... learning, as well as many different procedures that can be ... Michalski, 1980, for a survey). The most important of these ... learning are listed as follows.

a. Learning Control Information - This occurs when a prog... examine its own behavior in response to previous inputs an... results to discover where guessed correctly or incorrectly. ... then relate the results of this analysis to those par... of ... were responsible, it can modify such parts to improve its ... subsequent inputs. For example, a planner, during the pla... itself, can examine traces of its previous planning attempts ... when it investigated dead ends. If the planner uses rules ... sions as to which branch to investigate next, it can exam... to find those responsible for making bad choices and modify ... ingly. Automatic programming can also be viewed as the ... trol information to a logical specification of program. ... information can come from attempting to run the program ... and discovering empirically what must be added for succes... or by reasoning from first principles to determine what ... tion is needed.

b. Discovering Order in the World - In this case, real-world data are analyzed by the learning program to discern any regularities and then use them to make predictions or explain the data. There is a major dichotomy in the techniques used, depending on whether the data are expected to be almost "noiseless" or very "noisy." In the noiseless case, there are quite simple procedures for extracting the simplest generalization to explain the data, as for example, in finding a grammar that could have generated a given set of sentences or a taxonomy that gives the simplest classification of the given data. In the noisy case, however, the inductive search procedure must apply more complex statistical methods to decide if a particular "theory" of the data is supported by the data or is just a chance correlation.

Unfortunately, the amount of research in this basic area of AI is pitifully small compared with that in other areas. One reason for this is that learning is such a wide area that it is not a useful research strategy to try and develop a general-purpose learning system, but rather to concentrate on including learning in particular systems such as planners, natural-language systems, and automatic planners. Even with these limited aims, the amount of research is still small and fragmentary. This is unfortunate, since it is clear that the most direct road to truly intelligent flexible systems is to allow them to learn so they can improve themselves. Because of this potential and the need for autonomous systems to learn, this is one area of AI that NASA should seriously consider funding.

9. Summary

The major areas of AI basic research and applications have been reviewed. They are all different - each with its own particular relevance to NASA's aims and resources. In some areas, such as robot assembly in space, some of the basic research necessary has already been done, making it possible for NASA achieve useful progress in the near future. In other areas, such as fully autonomous remote-vehicle control, much research remains to be done.

This Page Intentionally Left Blank

VII CONCLUSIONS AND RECOMMENDATIONS

In the preceding discussion, we reviewed the current state of artificial-intelligence technology and considered a wide range of potential applications of AI in NASA. Some conclusions and recommendations with respect to an AI program for NASA can now be offered.

A. Conclusions

The number of prospective applications of AI in NASA is already very large - and growing apace. At present, NASA is developing only one or two applications, too few in view of the potential. Within the next five to ten years, the number of applications will increase dramatically, in a manner analogous to the proliferation of computer applications in NASA twenty years ago.

NASA needs AI to improve performance in existing tasks and to enable new tasks that would not be possible without AI.

The initial application being developed in NASA is an expert planning system, DEVISER (being developed at JPL), for mission planning. Although DEVISER does not represent the latest AI technology, it is nevertheless sound and a good base for additional development. Current AI planning technology can support enhancement of DEVISER as well as an expansion of its scope. In addition to DEVISER, a MITRE-developed planning system, KNOBS, is being adapted to crew activity planning by Goddard Space Flight Center. Other planning systems, such as SRI's SIPE, could also be adapted to NASA mission planning.

Mission planning has been identified by several past studies as a good prospect for application of AI technology, and such systems as DEVISER, KNOBS, or SIPE can be applied at Goddard Space Flight Center's Mission and Data Operations, Johnson Space Center's Operations Division, and at other centers.

An area related to mission planning - monitoring, diagnosis, and repair - has also been identified as important for the application of AI and in need of attention. Current AI technology could also support applications in this area.

NASA's AI requirements are not limited to a few AI techniques. NASA will need them all.

NASA is hardly alone, of course, in needing AI. AI technology will inevitably be developed over a broad spectrum of applications in the world outside NASA. However, NASA's priorities are different and its requirements for AI technology in space, such as for extremely autonomous robot systems, can justifiably be regarded as special.

In engineering, AI will be needed as an interface between users and complex data bases, as well as for planning, monitoring the progress of projects, etc. An important opportunity exists for AI in the development of the space station.

As regards management, AI systems will be used for improving the effectiveness of all the traditional functions, as "people amplifiers." In addition, AI planning systems can be especially helpful in planning for technological change.

NASA is becoming more and more dependent on computer data bases, in science, applications, operations, engineering, and many other areas of activity. AI techniques can be used to provide an intelligent interface

between users and collections of data bases that require different access languages and procedures (as illustrated in Figure 6). Also, future systems employing AI techniques will be more effective if current data-base development is done with such systems in mind.

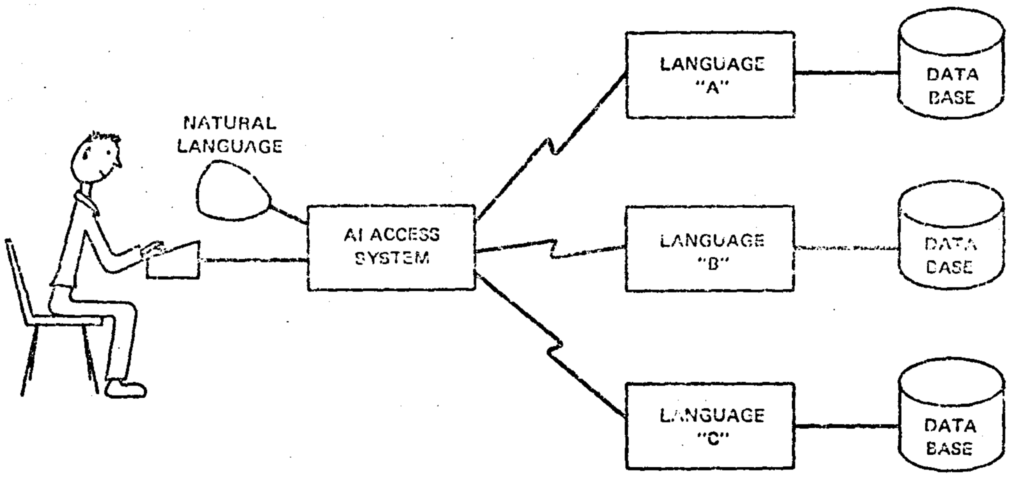

FIGURE 6 ACCESS TO A COLLECTION OF DATA BASES

Artificial intelligence will be essential for future NASA missions. The use of AI for fault diagnosis and repair in future aircraft and spacecraft has been identified as critical. Eventual space missions, such as a robot explorer, automated manufacturing, and self-replicating systems, will also require AI.

NASA will certainly benefit greatly from progress in automatic programming, as it will then take far less time, effort, and expertise to produce software.

B. Recommendations

1. Planning and Monitoring

DEVISER should be exploited immediately to serve as a nucleus system for one or two applications in mission planning, apart from the activity at JPL. DEVISER represents NASA's current experience in AI

systems, and has imparted as essential momentum that should not be lost. An effort to use DEVISER at a center such as Goddard will develop additional knowledge and experience within NASA and provide a base for expanding AI applications.

A project to use AI for monitoring spacecraft downlinks and diagnosing faults should be initiated along the lines already being considered in the Flight Control Division of Johnson Space Center.

2. Space Station

The space station is an opportunity, perhaps a necessity, for artificial intelligence. One study, conducted by MIT for Marshall Space Flight Center, concluded that AI may be indispensible for fault diagnosis and repair in planned space missions such as the space station. A knowledge-based AI system for the space station should be started now, beginning with conceptual design. The various data bases that will be created and used during the design, testing, operation, and maintenance of the space station should all become part of the knowledge base (an approach also advocated by the MIT study). This knowledge base could then accommodate many purposes, some of which would require AI techniques. The use of such techniques with the space-station knowledge base should begin as soon as possible. Early application of artificial intelligence will guide the development of the space-station data bases to make them more compatible with AI techniques, as well as providing an opportunity to test and verify the data bases. AI should be introduced into the design of the space station on a selective basis at first, but eventually with a sufficiently broad scope to encompass the entire knowledge base. All data base development for the space station, including IPAD data bases, should be planned for eventual operation with AI systems. The management of the space station project, in both its development and its implementation, should be supported by an expert planning system.

3. A NASA AI Research Group

Because of the importance of artificial intelligence in NASA's future, NASA needs an in-house research group. This group must work toward several objectives, but its basic function will be to establish NASA's membership in the AI research community. It should provide technical guidance for all the NASA applications of AI. It needs to understand and become an integral part of the AI research community — and, at the same time, of course, to understand and be a part of NASA. Initially concentrated at one center because of the scarcity of AI personnel, in-house AI capability should eventually become distributed in the NASA organization, just as computer technology has become so widely distributed in NASA during the past twenty years.

Membership in the AI research community would provide NASA with constantly updated knowledge about current AI technology and access to the research resources (including large amounts of software) shared by the community. The AI research community is presently quite small and primarily academic, consisting of approximately a thousand research professionals who maintain good mutual communication through the ARPANET, exchanges of visitors, seminars, society meetings, etc. Most importantly, the sophisticated, user-oriented, and highly effective AI software environment is at present cooperatively developed, maintained, and shared by the entire AI community.

The NASA AI research group would ensure the availability of pertinent current technology for NASA applications. As a related and very important benefit, it would also provide education and training for NASA personnel responsible for application development. Because experienced AI applications development personnel are scarce and difficult to secure, some current NASA people will need to be retrained.

An AI research group should be established at an appropriate NASA center, preferably near a major AI research community. Ames Research Center would appear to be an ideal choice. Five aspects of the recommended center are as follows:

- The AI research group should be modeled after a successful, established AI research facility, in both personnel and facilities (computers and software). It should grow to a total of twenty to thirty persons and expect to spend approximately two million dollars a year, with an investment of two to three million dollars in computers. As it will take several years to reach this level and, moreover, finding qualified personnel will not be easy, a start should obviously be made as soon as possible.

- An AI research facility similar to the one recommended here was described in a previous study for Goddard (Brown, October 1981). However, the facility recommended for Goddard was more applications-oriented and limited in scope than the research group recommended here. Figure 7 depicts a suggested arrangement for organization and staffing of the proposed research group, with a total of twenty persons assumed. The facility could be operated by contractor personnel and some fraction (but less than one-half) of the research staff could be comprised of visitors, either NASA personnel on loan for training or computer scientists from other AI research groups.

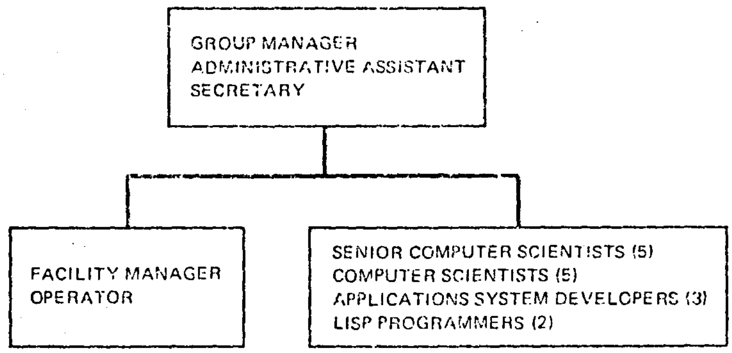

FIGURE 7 ORGANIZATION OF THE RECOMMENDED AI RESEARCH GROUP

- The center should share resources with the AI research community, making use of the ARPANET, visitor exchanges, guest lecturers, and such professional societies as the American Association for Artificial Intelligence.

- The center should take an aggressive role in identifying and assisting in the development of AI applications in NASA.

- Future applications of AI in NASA must be consistent with overall NASA programs and objectives; this requires

systematic guidance by management and an overall understanding of NASA on the part of the AI research group. Consequently, a good management structure extending upwards to the topmost echelons of the NASA organization needs to support, and be supported by, the AI research group.

Figure 8 illustrates the position of the NASA AI research group in relation to the AI research community at large, to applications of AI in NASA, and to NASA management.

The application of AI systems is a new endeavor. Consequently, experience in developing AI applications outside the research laboratory is practically nonexistent. In general, AI research personnel are not really interested in developing practical applications, so that future professionals who can devote themselves to such development will, for the most part, have to come from elsewhere. NASA will need to create its own personnel. It already has individuals with experience in system development, but without knowledge of artificial intelligence. The obvious solution is to start by utilizing and retraining these people. The proposed research group would facilitate this effort by having NASA system development personnel assist in its AI research activities.

All of the recommendations presented here have far-reaching implications for decades to come. Each calls, however, for immediate action: (1) for mission planning, applications development at one or more centers such as Goddard and JSC; (2) for the space station, a commitment to design principles that accommodate AI techniques; (3) for AI research in NASA, an appropriately organized and supported group at Ames Research Center.

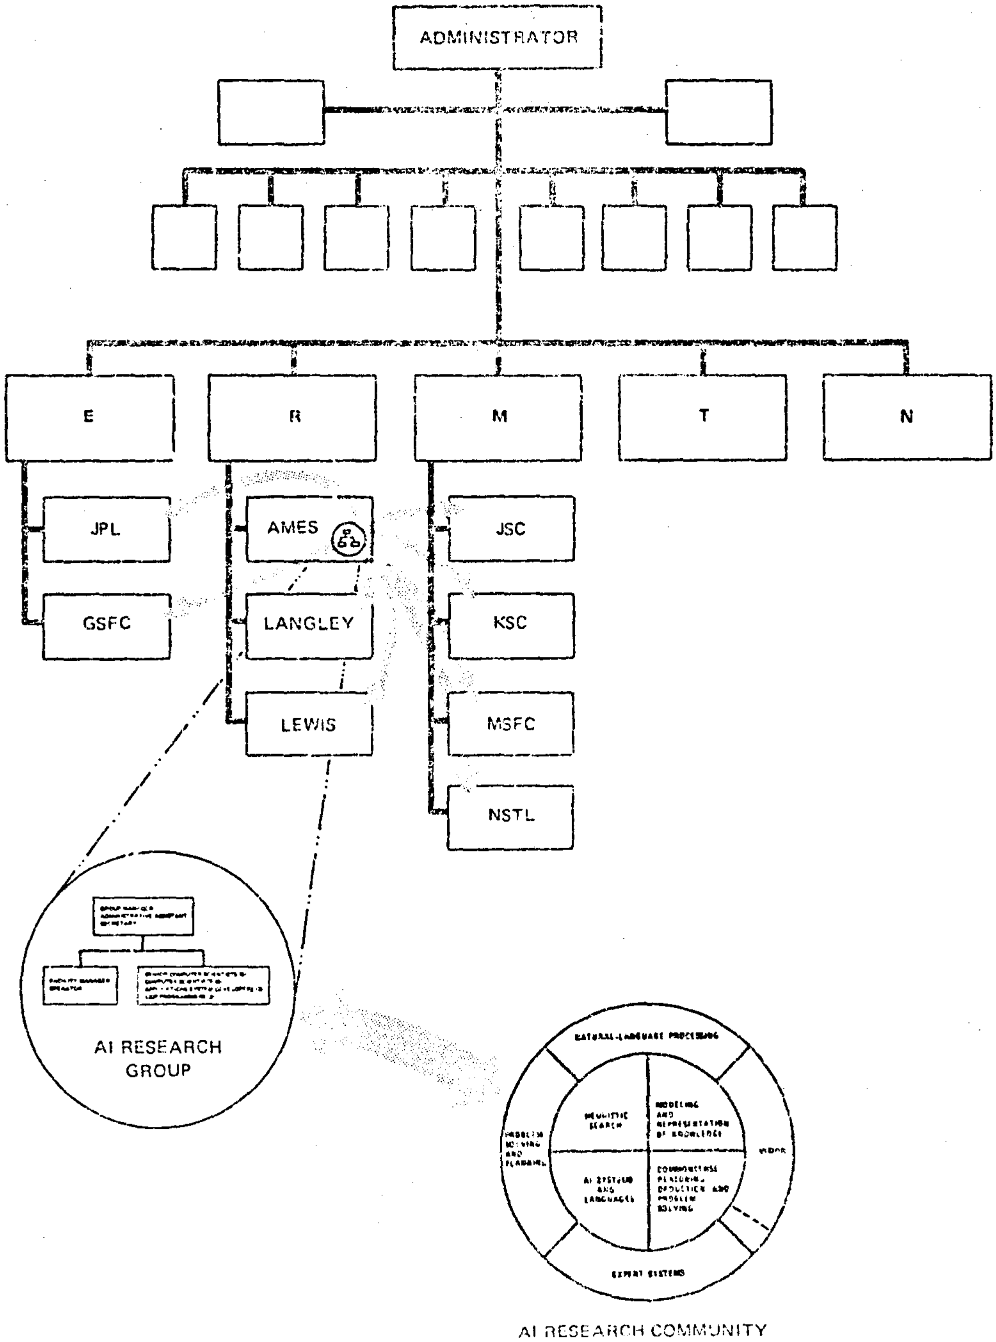

FIGURE 8 THE NASA AI RESEARCH GROUP IN RELATION TO OTHER ORGANIZATIONS

BIBLIOGRAPHY

Albus, James A., People's Capitalism: The Economics of the Robot Revolution (New World Books, College Park, Maryland, 1976).

Barr, Avron, and Edward A. Feigenbaum, The Handbook of Artificial Intelligence, Vol. 2 (William Kaufmann, Inc., Los Altos, California, 1982).

Bekey, Ivan, and Harris Mayer, "1980-2000: Raising Our Sights for Advanced Space Systems," Astronautics & Aeronautics, Vol. 14, No. 7/8, pp. 34-63 (July/August 1976).

Bekey, Ivan, and J. E. Naugle, "Just Over the Horizon in Space," Astronautics & Aeronautics, Vol. 18, No. 5, pp. 64-76 (May 1980).

Brooks, Rodney A., "Symbolic Reasoning Among 3-D Models and 2-D Images", Artificial Intelligence, Vol. 17, Nos. 1-3, pp. 285-348 (August 1981).

Brown, David R., "Applications of Artificial Intelligence in Space Tracking and Data Systems," Task 1 Report, Contract NAS5-26358, SRI Project 2203, SRI International (April 1981).

Brown, David R., "Recommendations for an AI Research Facility at NASA/GSFC," Task 2 Report, Contract NAS5-26358, SRI Project 2203, SRI International (October 1981).

Buchanan, B. G., and E. A. Feigenbaum, "DENDRAL and META-DENDRAL, Their Applications Dimension," Artificial Intelligence, Vol. 11, Nos. 1,2, pp. 5-24 (August 1978).

"The CETI Program," Soviet Astronomy, Vol. 18, No. 5, pp. 669-675 (March-April 1975), a translation from Astronomicheskii Zhurnal, Vol. 51, No. 5 (September-October 1974), pp. 1125-1132.

Duda, Richard O., "The PROSPECTOR System For Mineral Exploration", Final Report, Contract No. 14-08-0001-17296, SRI Project 8172, SRI International, Menlo Park, California (April 1980).

Feigenbaum, Edward A. and Pamela McCorduck, The Fifth Generation (Addison-Wesley Publishing Company, Reading, Massachusetts, 1983).

Fikes, Richard E., Peter E. Hart, and Nils J. Nilsson, "Learning and Executing Generalized Robot Plans", Artificial Intelligence, Vol. 3, No. 4, pp. 251-288, (Winter 1972).

"A Forecast of Space Technology: 1980-2999," NASA SP-387, Scientific and Technical Information Office, National Aeronautics and Space Administration, Washington, D.C. (1976).

Freitas, Robert A., and Patricia A. Carlson, eds., "Computer Science: Key to a Space Program Renaissance," Technical Report 1168, University of Maryland, College Park, Maryland (January 15, 1982).

Freitas, Robert A., and William P. Gilbreath, eds., "Advanced Automation for Space Missions," NASA CP-2255, Scientific and Technical Information Branch, National Aeronautics and Space Administration (1982).

Fulton, Robert E., "CAD/CAM Approach to Improving Industry Productivity Gathers Momentum," Astronautics & Aeronautics, Vol. 20, No. 2, pp. 64-70 (February 1982).

√ Gevarter, William B., "An Overview of Robotics," NBSIR 82-2479, National Bureau of Standards, Washington, D.C. (March 1982).

Gevarter, William B., "An Overview of Expert Systems," NBSIR 82-2505, National Bureau of Standards, Washington, D.C. (May 1982).

√ Gevarter, William B., "An Overview of Computer Vision," NBSIR 82-2582, National Bureau of Standards, Washington, D.C. (September 1982).

Heer, Ewald, "New Luster for Space Robots and Automation," Astronautics & Aeronautics, Vol. 16, No. 9, pp. 48-60 (September 1978).

Heer, Ewald, "Report of the JPL Advocacy Group for Autonomous Systems Technology Research and Development," 715-128, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California (June 1981).

Johnson, Richard D., and Charles Holbrow, eds., "Space Settlements: A Design Study," SP-413, Scientific and Technical Information Office, National Aeronautics and Space Administration, Washington, D.C. (1977).

"Machine Intelligence and Robotics: Report of the NASA Study Group," Executive Summary, 730-51, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California (September 1979).

"Machine Intelligence and Robotics: Report of the NASA Study Group," Final Report, 715-32, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California (March 1980).

McCarthy, J., P. W. Abrahams, D. J. Edwards, T. P. Hart, and M. I. Levin, LISP 1.5 Programmer's Manual, MIT Press, Cambridge, Mass. (1965)

McDermott, John, "R1: A Rule-Based Configurer of Computer Systems," Artificial Intelligence, Vol. 19, No. 1, pp. 39-88 (September 1982).

McReynolds, S. R., "A Benefit and Role Assessment of Advanced Automation for NASA," JPL 730-13, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California (14 April 1978).

Meintel, Alfred J., and Ronald L. Larsen, "NASA Research in Teleoperation and Robotics," Robotics and Industrial Inspection, SPIE Proceedings, Vol. 360, pp. 22-30 (Society of Photo-Optical Instrumentation Engineers, Bellingham, Washington, 1983).

Michalski, Ryszard S., "Pattern Recognition as Rule-Guided Inductive Inference", IEEE Transactions on Pattern Analysis and Machine Learning, Vol. PAMI-2 No. 4, pp. 349-361 (July 1980).

Miller, Rene H., Marvin L. Minsky, and David B. S. Smith, "Space Applications of Automation, Robotics and Machine Intelligence Systems (ARAMIS)," Volume 1: Executive Summary, Phase 1, Final Report, NASA CR-162079, Space Systems Laboratory and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts (August 1982).

Morrison, Philip, John Billingham, and John Wolf, eds., "Search for Extraterrestrial Intelligence," SP-419, Scientific and Technical Information Office, National Aeronautics and Space Administration, Washington, D.C. (1977).

"NASA Space Systems Technology Model," Code RS, Office of Aeronautics and Space Technology, National Aeronautics and Space Administration (May 1980).

"National Aeronautics and Space Act of 1958, as Amended, and Related Legislation," prepared at the request of the Committee on Commerce, Science, and Transportation, United States Senate, 95th Congress, 2nd Session (December 1978).

Nii, H. Penny, Edward A. Feigenbaum, John J. Anton, and A. J. Rockmore, "Signal-to-Symbol Transformation: HASP/SIAP Case Study," AI Magazine, Vol. 3, No. 2, pp. 23-35 (Spring 1982).

"Outlook for Space: Report to the NASA Administrator by the Outlook for Space Study Group," NASA SP-386, Scientific and Technical Information Office, National Aeronautics and Space Administration (1976).

Pearl, Judea, "Heuristic Search Theory: Survey of Recent Results," Proceedings of the Seventh International Joint Conference on Artificial Intelligence (IJCAI-81), Vancouver, British Columbia, Canada, 25-28 August 1980, pp. 554-562 (American Association for Artificial Intelligence, Menlo Park, California, 1981).

Shortliffe, E. H., Computer-Based Medical Consultations: MYCIN (Elsevier, New York, 1976).

Sussman, Gerald Jay, "Electrical Design: A Problem for Artificial Intelligence Research," Proceedings, Fifth International Joint Conference on Artificial Intelligence, Massachusetts Institute of Technology, Cambridge, Massachusetts, 22-25 August 1977, pp. 894-900.

Vere, Steven A., "Planning in Time: Windows and Durations for Activities and Goals," Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California (November 1981).

von Tiesenhausen, Georg, and Wesley A. Darbro, "Self-Replicating Systems - A Systems Engineering Approach," NASA TM-78304, Marshall Space Flight Center, Alabama (July 1980).

Weiss, Sholom M. and Casimir A. Kulikowski, "EXPERT: A System for Developing Consultation Models," Proceedings, Sixth International Joint Conference on Artificial Intelligence, Tokyo, Japan, 20-24 August 1979, pp. 942-947.

END DATE

JAN. 19, 1984

End of Document